

# Wright meets Markowitz: How standard portfolio theory changes when assets are technologies following experience curves

Rupert Way <sup>\*a, b</sup>, François Lafond <sup>†a, b, c</sup>, Fabrizio Lillo <sup>‡d, e</sup>, Valentyn  
Panchenko <sup>§f</sup> and J. Doyne Farmer <sup>¶a, g, h, i</sup>

<sup>a</sup>*Institute for New Economic Thinking at the Oxford Martin School,  
University of Oxford, Oxford OX2 6ED, UK*

<sup>b</sup>*Smith School of Enterprise and the Environment, University of Oxford,  
Oxford, OX1 3QY*

<sup>c</sup>*Oxford Martin School Programme on Technological and Economic  
Change, University of Oxford, OX1 3BD*

<sup>d</sup>*Scuola Normale Superiore, Piazza dei Cavalieri 7, 56126 Pisa, Italy*

<sup>e</sup>*CADS, Center for Analysis, Decisions, and Society, Human Technopole,  
Milano, 20156, Italy*

<sup>f</sup>*School of Economics, UNSW Business School, Sydney NSW 2052,  
Australia*

<sup>g</sup>*Mathematical Institute, University of Oxford, Oxford OX1 3LP, UK*

<sup>h</sup>*Department of Computer Science, University of Oxford, Oxford OX1  
3QD, UK*

<sup>i</sup>*Santa-Fe Institute, Santa Fe, NM 87501, USA*

August 26, 2018

---

\*rupert.way@smithschool.ox.ac.uk

†francois.lafond@inet.ox.ac.uk

‡fabrizio.lillo@sns.it

§v.panchenko@unsw.edu.au

¶doyne.farmer@inet.ox.ac.uk

## Abstract

We consider how to optimally allocate investments in a portfolio of competing technologies using the standard mean-variance framework of portfolio theory. We assume that technologies follow the empirically observed relationship known as Wright’s law, also called a “learning curve” or “experience curve”, which postulates that costs drop as cumulative production increases. This introduces a positive feedback between cost and investment that complicates the portfolio problem, leading to multiple local optima, and causing a trade-off between concentrating investments in one project to spur rapid progress vs. diversifying over many projects to hedge against failure. We study the two-technology case and characterize the optimal diversification in terms of progress rates, variability, initial costs, initial experience, risk aversion, discount rate and total demand. The efficient frontier framework is used to visualize technology portfolios and show how feedback results in nonlinear distortions of the feasible set. For the two-period case, in which learning and uncertainty interact with discounting, we compare different scenarios and find that the discount rate plays a critical role.

# 1 Introduction

There is a fundamental trade-off, encountered throughout life, between investing enough effort in any one activity to make rapid progress, and diversifying effort over many projects simultaneously to hedge against failure. On the one hand, by focusing on a single task one can quickly accumulate experience, become an expert, and reap rewards more efficiently. But on the other, unforeseen circumstances can impede progress or make the rewards less valuable, so it may be wise to maintain progress on several fronts at once, even if individually slower. This brings to mind the familiar adage “don’t put all your eggs in one basket”, and at first glance appears very similar to the question of how diversified a portfolio of financial investments should be. However there is a key difference between the financial portfolio setting and the type of problem considered in this paper, which is that here *learning* is involved: the more effort we invest in one area, the more *effective* that effort becomes — so we may want to put all our eggs in one basket after all.

The dilemma is ubiquitous, and is understood intuitively by us all as we learn new skills, engage in new projects, and attempt to plan for the future. For example, consider trying to decide how many courses to take in university; or how many languages, musical instruments, sports, or web application frameworks to learn. Focusing on one, or just a few, allows us to gain expertise and reach a more rewarding phase of activity sooner. Or at the organisational level, firms and governments must decide how many, and which, strategic and technological capabilities to develop. We present a simple model for understanding this trade-off, and show how it is related to the optimal diversification problem for financial assets.

The reason this decision framework is of particular interest is that, despite its simplicity, it shares several important features with the question of how to allocate investments among competing technologies. This is because often, in the long run, scientific advances and knowledge gains mean that performance-weighted technology investment costs decrease as cumulative deployment increases. Put simply, in such cases the more we invest in a technology (whether at the R&D, deployment or any other stage) the more effective the technology becomes at delivering the same output, so future investment costs are lower, per unit of output<sup>1</sup>. Hence, in order to achieve certain long term technological goals, understanding the correct allocation of investments among available substitute technologies is vital. The specific question we have in the back of our minds is how to allocate funding over potential clean energy technologies to accelerate the transition to a net zero carbon economy — should we invest in solar photovoltaics, or offshore wind, or next generation nuclear, or carbon capture and storage, or a little bit in each?

Despite this high-level motivation, here we focus in on a very simple conceptual model representing the underlying trade-off. The key assumption we make is that increased cumulative investment in a technology *leads to* reduced investment costs (but with some degree of uncertainty). In reality this causal mechanism is not so straightforward, and there are many other complicating factors, such as correlations between projects and spillovers (incoming and outgoing) of various kinds. However, while stating the caveats clearly, we set aside these issues for now and just focus on the core problem, which is to find the optimal risk-averse investment in competing technologies following experience curves. This setting brings together the specialization incentives of the learning

---

<sup>1</sup>Though note that many technologies do not exhibit such decreasing costs at all, and also that the effect (when it exists) is far less evident in mature technologies since so much experience has already accrued.

curve model with the diversification incentives of modern portfolio theory, allowing us to characterize the optimal solution to the trade-off between diversifying and specializing.

Our approach is to consider multiple independent technologies (two), increasing returns to investment (through experience curves), uncertainty (cost is a stochastic process), and a risk-averse decision maker (who minimizes a mean-variance value function). Since the two technologies follow stochastic processes, diversification tends to reduce the risks, but at the same time increasing returns tend to favour specialization. Investing in one option drives down its marginal cost, making it more and more attractive, but *ex-ante* uncertainty in future benefits from learning suggests that diversifying can limit the risk of over-investing in a technology that eventually shows a poor performance. We characterize optimal investment as a function of learning characteristics (rate and uncertainty of learning), initial conditions (cost competitiveness and accumulated experience), risk aversion, discount rate and the level of demand. We focus primarily on the one-period investment decision, but also consider the extension to two periods.

In classical (Markowitz) portfolio theory, the optimal allocation of investments is unique. In general, the further a portfolio is from the optimum, the worse is its value. In contrast, when the positive feedback of endogenous technological progress is strong, it is better to invest mostly in *either* of the two options than to split investment more evenly. Except in some knife-edge cases, one of the two specialized portfolios is better than the other, but which one is best depends on the parameters. As a result, a small change in one of the parameters can result in the optimal portfolio being completely different.

In general we characterize three different regimes. In the first regime, one technology is so much better than the other that it dominates the portfolio entirely. This happens either because there is no risk aversion (so we revert to the classic deterministic learning curve winner-takes-all scenario) or because the relative advantage of one technology in terms of initial conditions or speed and uncertainty of learning is very strong. In the second regime, the optimal portfolio features unambiguous diversification, that is, the objective function has a unique optimum corresponding to a balanced mix of technologies. In the third regime, there are two local optima of similar value, corresponding to quite different investment policies. As parameters change, the transition of the global optimum from one local optimum to the other is abrupt. In other words, in this critical region, a small change in a parameter causes a large change in the optimal policy.

We show what this finding implies for the theory of path-dependence and lock-in, and characterize lock-in as a situation where investing in a fast-learning technology is not currently optimal, but it *would be* if a higher level of demand existed. Intuitively, whether or not one should attempt to bring a technology down its learning curve depends on the size of the market. For some parameter values though, the transition is sharp – there exists a critical level of demand below which investment in the fast-learning technology is limited, and above which it becomes dominant (i.e. the global optimum switches from one local optimum to the other).

We show analytically that a Markowitz-like case may be recovered in two different ways. First, when there is no learning increasing returns are absent and it becomes highly unlikely that one would want to specialize entirely. (Specialization is still possible, but only because one technology is currently much better than the other, *not* because investing in it makes it better.) Second, when future demand is very small, compared to the current level, the potential for learning is insignificant and therefore investment is never enough for increasing returns to really matter.

Our results relate to several different branches of literature. First of all we are moti-

vated by the optimal energy systems, energy transition and climate change literatures. It is now clear that to avoid a rise in temperatures that would have “dangerous” effects we must attain net zero carbon emissions, and one of the main factors involved in this is a switch to a clean energy system. However, dirty energy technologies are currently considered to be both cheap and convenient (in some ways due to legacy energy system design and infrastructure, and societal embeddedness), whereas alternatives are still expensive, even though their cost is falling, sometimes very fast. In this context the questions arise of what costs will be in the future, which decisions affect these costs, and what is the best investment or tax/subsidy policy. Energy systems are highly complex<sup>2</sup> and energy experts generally rely on highly detailed models of the energy production and consumption mix. To include endogenous technological change in these models, a simple solution which has been widely adopted (and criticized) is that of experience curves (Gritsevskiy & Nakićenović 2000, Barreto & Kypreos 2004, Alberth & Hope 2007, Criqui et al. 2015, Webster et al. 2015). However, these models end up being very complex so that optimal policies are very hard to determine and understand. Numerical methods have to be used, and it is not always the case that the global optimum is found. Analytical approaches for these complex models generally have to assume a deterministic setting, so that the increasing returns induced by the learning curve lead to full specialization, see for instance Wagner (2014). In this paper we only wish to understand the fundamental trade-off involved in technology investment: diversification against specialization, and the risk of lock-in in systems with path dependent, self-reinforcing dynamics. Therefore, we do not attempt to provide a realistic model of the energy system or a direct empirical application of our results, and instead focus on a theoretical contribution at the intersection of the learning-by-doing and portfolio literatures.

To model technological progress, we use a very specific parametric model. Technological progress is not perfectly predictable, but in many detailed empirical cases it has been found that unit costs tend to decrease by a constant percentage every time cumulative production doubles. Subject to some uncertainty about future shocks, the cost of a technology follows an experience curve which is technology-specific. This relationship between unit cost and cumulative investment has been observed for a long time (Wright 1936, Alchian 1963, Thompson 2012) and is generally explained by the fact that during production learning-by-doing takes place<sup>3</sup>. Starting with Arrow (1962), a large literature has developed to analyse the consequences of this relationship for pricing and output decisions (Rosen 1972, Spence 1981, Mazzola & McCardle 1997). Learning-by-doing decreases marginal cost, which gives an advantage to size and may encourage predatory pricing (Cabral & Riordan 1994) or legitimize the protection of infant industry from international competition (Dasgupta & Stiglitz 1988). When one considers a single firm operating a single technology, learning-by-doing generates irreversibilities and creates an incentive to delay investment. The optimal investment dynamics can be characterized

---

<sup>2</sup>Each energy source has specific infrastructure building time (e.g. nuclear takes a long time), can be intermittent or not (e.g. solar energy is not produced at night), has specific transport, storage and safety conditions, etc.

<sup>3</sup>While learning-by-doing often refers to labour force or organisational learning in particular, two other related and noteworthy sources of increasing returns include economies of scale, which depend only on current (not cumulative) production levels, and network externalities, which depend on the number of consumers or other producers joining or using the same network or technology. The key feature of experience curves, however, is that performance increases depend on the growth of total experience (cumulative production), not on the growth of production. We think of experience curves as capturing all experience-related effects, including, but not limited to learning-by-doing.

using the theory of real options (Brueckner & Raymon 1983, Majd & Pindyck 1989, Della Seta et al. 2012). In general, the literature does not study multiple technologies at the same time; and when it does, for instance when characterizing the social optimum for a multi-firm sector, it is generally in the absence of uncertainty. Given our motivation to understand optimal investment in energy technologies, which are very diverse and uncertain, we turn to another branch of literature which has dealt in detail with investment in multiple uncertain assets, that is modern portfolio theory (Markowitz 1952).

Modern portfolio theory considers a risk-averse decision maker who wishes to invest in financial assets. The key result of portfolio theory is that there exists an optimal way of combining assets in a portfolio such that expected returns are maximized, conditional on a given level of risk (or that risk is minimized, conditional on a given level of expected returns). We argue that this idea is well suited for thinking about technology investment, and we borrow from portfolio theory the mean-variance value function (in our case, both expected costs and variance of the portfolio have to be minimized). For simplicity, however, we generally assume that technologies are uncorrelated. As opposed to a “learning curve technology”, a key property of a financial asset is that investing in it does not change its value, although there are some important exceptions<sup>4</sup>. We recover a classical portfolio setup when the learning parameter is zero or when the total market size is very small compared to the initial production.

Besides our general motivation (energy systems) and the two major ingredients of our model (experience curves and mean-variance portfolio theory), our setup relates to a large literature dealing with optimal control of stochastic processes, which goes well beyond economics and operation research. Of more direct interest are the applications to technology, R&D and innovation problems, where the questions of increasing returns and lock-in are more salient. When investing in an option makes it better and better, history matters. Atkinson & Stiglitz (1969) already pointed out that localized technological progress, an important source of which is learning-by-doing, would justify investing in a technology that is not yet the cheapest. In the technology choice literature, it is well known that increasing returns and uncertainty may result in situations where poor technological options dominate (David 1985). In a model of two competing standards operating under network externalities, Arthur (1989) showed that if chance favors an intrinsically worse option early on, this option’s accumulated experience gives it an edge for obtaining the marginal consumer. As this advantage accumulates, it may forever exceed the benefits from switching to the intrinsically better option. In this context a policy maker is interested in a policy that optimally explores the merit of different options before making a final choice. Cowan (1991) characterized a social planner’s optimal decision in a two arm bandit framework, where there is a choice between one of two technologies at every period. In this model, there exists an optimal policy known as the Gittins index, but according to this policy eventually a single technology will be chosen. Thus early bad luck may induce the social planner to lock in the wrong technology. While

---

<sup>4</sup>One is the situation in which market impact is considered. Market impact acknowledges that trading large quantities simply violates the atomicity assumption, so that one’s choice of quantities demanded or supplied affects the price. In this case, this is a *negative* feedback and the literature has focused on finding optimal liquidation strategies (Almgren & Chriss 2001, He & Mamaysky 2005). Another situation in which financial portfolios incorporate feedback effects is when learning about an asset is taken into account. An investor who is familiar with a particular asset makes more precise estimates of expected returns, so that this asset is relatively more valuable than other assets (Boyle et al. 2012). In turn, holding a lot of a particular asset makes information acquisition about that asset more valuable, which can generate a positive feedback that encourages specialization (Van Nieuwerburgh & Veldkamp 2010).

this theoretical literature often refers to learning-by-doing<sup>5</sup>, it attempts to model other forms of increasing returns all at once and therefore does not model more explicitly how cost decreases with investment. Zeppini (2015) considered learning curves for clean and dirty technologies in a discrete choice framework, with social interactions as an additional source of increasing returns to adoption and lock-in. He found that policies inducing the clean technology to progress down its learning curve faster have greater potential to induce smooth technological transitions, as opposed to traditional policies such as a pollution tax which can work only by being large enough to induce an equilibrium shift. Finally, another branch of literature has contrasted the benefits of increasing returns against the benefits of technological diversity by assuming that further technological progress takes place through recombination. This implies that there is some value in giving up on increasing returns from specialization and keeping a range of diverse technologies available for further re-combination (Van den Bergh 2008, Zeppini & Van den Bergh 2013).

The paper is organized as follows. Section 2 defines the stochastic process for the experience curves and the optimization problem in the one-period case, and shows how it relates to Markowitz portfolios. Section 3 presents the main results of the optimization and shows under which conditions diversification is optimal. It also analyzes in detail the objective function by characterizing how the number and nature of optima changes with underlying parameter values, and studies the effect of total demand. Section 4 returns to the comparison of financial and technology portfolios and shows how the efficient frontier changes when technologies are introduced. Section 5 establishes conditions to escape lock-in by studying the case where a mature, cheap but slow-learning technology dominates the market but faces competition from a young, expensive but fast-learning challenger. Section 6 introduces the multi-period model and explores how discounting interacts with risk aversion and learning in a two-period setting. Finally, Section 7 concludes.

## 2 One-period model

Consider the development of a single technology over one time period. The unit cost of the technology at time  $t$  is  $c_t$  (measured in  $\$/unit$ ), and its cumulative production<sup>6</sup> (measured in *units*) is  $z_t$ . Let  $t = 0$  be the present time and  $t = 1$  be some given future time. The current unit cost is  $c_0$  and current cumulative production is  $z_0$ . Production during the period is  $q$ , and the cumulative production at  $t = 1$  is  $z_1 = z_0 + q$ . We first present the stochastic model for a single technology then consider a portfolio of two such technologies.

### 2.1 Wright’s law

The standard form of the experience curve is

$$c_t \propto z_t^{-\alpha}, \quad (1)$$

where the constant  $\alpha$  is the experience exponent (or Wright exponent) for this technology. This leads to two related concepts often used in the literature: the “progress ratio” is

---

<sup>5</sup>When increasing returns are from the consumer side, typically as in Arthur (1989), they are generally motivated as learning-by-using following Rosenberg (1982).

<sup>6</sup>We use the terms investment and production interchangeably throughout the one-period model presentation. Generally the literature considers production, although the original paper by Arrow (1962) used investment. Here we are looking only one step ahead so this is not an important difference.

defined as the relative cost level seen after each doubling of cumulative production,  $PR = 2^{-\alpha}$ , while the “learning rate” is defined as the relative cost reduction seen after each such doubling,  $LR = 1 - 2^{-\alpha}$ . Dutton & Thomas (1984) report learning rates from different studies and find that the vast majority lie between 5% and 40%, corresponding to values of  $\alpha$  lying approximately within the range (0.07, 0.7). However, commodities such as minerals and fossil fuels mostly have  $\alpha \approx 0$  since they do not exhibit a significant cost decrease over the long run (Newbold et al. 2005, McNerney et al. 2011). The power law relationship between cost and cumulative production was first noted by Wright (1936) in the context of the production of airplanes, so we call it Wright’s law. Since then it has been found to describe the available evidence for a number of technologies fairly well (Nagy et al. 2013). In contrast to a large part of the theoretical literature on experience curves, which deals only with the deterministic form, we model uncertainty explicitly. To do this we make the future cost stochastic by assuming additive noise  $\eta$  on the log-first-difference version of Eq. (1):

$$\log(c_1) - \log(c_0) = -\alpha [\log(z_1) - \log(z_0)] + \eta. \quad (2)$$

This equation models a situation where, over the course of one period, an underlying linear trend in log-log space advances according to Wright’s law, but then is hit by a random shock. It is one of the simplest possible ways of incorporating uncertainty in the experience curve model, chosen here specifically for its clarity and simplicity<sup>7</sup>. The cost of production at  $t = 1$ , interpreted as the average (or constant) within-period cost, is then given by

$$c_1 = c_0 \left( \frac{z_0}{z_1} \right)^\alpha e^\eta = c_0 \left( \frac{z_0}{z_0 + q} \right)^\alpha e^\eta. \quad (3)$$

So there is a distribution of possible future costs  $c_1$ , and Eq. (3) shows clearly how it depends on: *i*) the current state,  $c_0$ ,  $z_0$ , of the technology<sup>8</sup>, *ii*) the technology’s experience exponent  $\alpha$ , *iii*) the choice of production  $q$  over the period, and *iv*) the noise distribution  $\eta$ .

Next, we suppose that the shock is normally distributed with mean zero<sup>9</sup> and variance  $\sigma^2$ ,  $\eta \sim \mathcal{N}(0, \sigma^2)$ . This noise model is known to be a reasonable assumption for financial assets with lognormal returns, but some justification is required when considering technologies. Lafond et al. (2018) found that this model gave a reasonably good fit to data on 51 technology time series, in the sense of predicting theoretical forecast errors in line with realised forecast errors, although their preferred model allows for autocorrelation.

Thus cost is log-normally distributed, and by standard log-normal properties its expectation and variance are given by

$$\mathbb{E}[c_1] = c_0 \left( \frac{z_0}{z_0 + q} \right)^\alpha e^{\sigma^2/2}, \quad (4)$$

$$\text{Var}(c_1) = c_0^2 \left( \frac{z_0}{z_0 + q} \right)^{2\alpha} e^{\sigma^2} (e^{\sigma^2} - 1). \quad (5)$$

---

<sup>7</sup>Another way would be to make the learning rate  $\alpha$  stochastic, instead of the cost. Mazzola & McCardle (1996) considered how a Bayesian learner benefits from more production not only by decreasing costs, but also by improved estimates of the learning parameter.

<sup>8</sup>Note that it is the presence of  $z_0$  here that distinguishes between learning effects and increasing returns to scale.

<sup>9</sup>Nonzero mean noise is discussed in Section 6.



These two properties of the stochastic experience curve, specified uniquely by the four parameters  $c_0, z_0, \alpha, \sigma$ , will now be used to construct the portfolio model.

## 2.2 The optimization

Consider two independent technologies,  $A$  and  $B$ , each evolving according to the form of Wright's law proposed above, with their own technology-specific parameters. We label variables and parameters with superscripts (e.g.  $q^A, c_0^A, z_0^A, \alpha^A, \sigma^A$ ). Suppose the technologies are perfect substitutes<sup>10</sup> and that there is a fixed, exogenous demand  $K$ , which must be satisfied exactly by some combination of production of the two technologies<sup>11</sup>, i.e. there is a production constraint  $K = q^A + q^B$ . Production is non-negative, so  $q^A, q^B \in [0, K]$ , and choosing  $q^A$  also determines  $q^B = K - q^A$ . We use  $q^A$  as the control variable in the following optimization and present results in terms of the share of total production in technology  $A$ ,  $q^A/K$ . Let the total cost of production during the period be  $V(q^A)$ . This is just the sum of unit costs times units produced

$$V(q^A) = \sum_{i=A,B} c_1^i q^i, \quad (6)$$

where stochastic costs  $c_1^i$  depend nonlinearly on productions  $q^i$ , as in Eq. (3). Thus for a fixed, known set of technology parameters  $\{c_0^i, z_0^i, \alpha^i, \sigma^i\}_{i=A,B}$  and total demand  $K$ , each choice of production  $q^A$  maps to a distribution of total costs  $V$ . The tools for addressing this type of problem are well developed, see for example Krey & Riahi (2013). The goal here is to understand how the parameters and the choice of production together generate the total system cost distribution, from which an optimal production portfolio may be identified. We perform a mean-variance analysis on  $V$  because it is simple, intuitive and illustrates clearly the key features of the system<sup>12</sup>. Let  $\lambda \geq 0$  be a risk aversion parameter and  $f$  be the mean-variance objective function. The optimization problem is then

$$\underset{q^A}{\text{minimize:}} \quad f(q^A) = \mathbb{E} [V(q^A)] + \lambda \text{Var} (V(q^A)) \quad (7)$$

$$\text{subject to: } q^A \in [0, K].$$

The aim therefore is to find the production mix which, while meeting the production constraint, minimizes the expected total cost of production, plus an additional term characterizing the spread of the distribution of possible outcomes. The risk aversion parameter  $\lambda$  scales the contribution of the variance term in  $f$ , reflecting the extent to

---

<sup>10</sup>While the perfect substitutability assumption is essential in this model, in reality technologies are often not continuously varying substitutes, and it may not be possible to adopt just a bit of several different technologies. Indeed, many technology adoption decisions are entirely binary, such as the choice of firm-wide software systems. This is a limitation of the model, and the domain of application should therefore be chosen carefully.

<sup>11</sup>Since demand and total production are assumed equal throughout we use the terms interchangeably. It is assumed that the demand is inelastic and prices are determined competitively (as is typical in energy markets). Under these conditions cost minimization is equivalent to profit maximization.

<sup>12</sup>Since empirical technology cost noise shocks are found to fit a lognormal distribution fairly well, as discussed previously, standard results from the finance literature apply here. In particular, use of the mean-variance decision framework in the one-period setting is justified as it provides a good approximation to all commonly used utility functions (Pulley (1981), Kroll et al. (1984)). However, the choice of utility function in a multi-period setting (as we consider in Section 6) is much more subtle, and a different objective function may be preferable.

which the decision maker prefers to minimize exposure to cost uncertainty. In the risk-neutral case ( $\lambda = 0$ ) the variance term has zero weight so the optimization just discovers the production mix with lowest expected total cost (in this case just a single technology). Conversely, in the high risk aversion case ( $\lambda \gg 1$ ) the second term in  $f$  dominates the first and so the optimization discovers the production mix with lowest total cost uncertainty, regardless of its expectation. In the intermediate regime both terms play a significant role in determining the outcome of the optimization. Using Eqs. (4), (5) and (6) the objective function in problem (7) may be written explicitly as

$$f(q^A) = \sum_{i=A,B} c_0^i \left( \frac{z_0^i}{z_0^i + q^i} \right)^{\alpha^i} e^{(\sigma^i)^2/2} q^i + \lambda \left( c_0^i \left( \frac{z_0^i}{z_0^i + q^i} \right)^{\alpha^i} q^i \right)^2 e^{(\sigma^i)^2} (e^{(\sigma^i)^2} - 1). \quad (8)$$

Thus  $f$  is just the sum of one cost-expectation-based component and one cost-variance-based component for each technology; covariance terms are zero due to the technology independence assumption (i.e.  $\eta^A$  and  $\eta^B$  are uncorrelated). (The case of correlated noise is considered in Section 3.5.)

This is a non-convex optimization problem so it may have more than one local minimum. Since there is only one free variable though,  $q^A$ , it is relatively quick to solve by brute force optimization. Denote the optimum by  $q_*^A$ .

Despite the simplicity of the model, the scope for understanding its behaviour via standard analytical techniques is rather limited. This is because the product terms  $(z_0^i + q^i)^{-\alpha^i} q^i$  in the objective function mean that differentiation of  $f$  just generates more and more similar product terms, which makes closed-form expressions for optima or other system properties only possible in a few restricted cases. Most of our results and analysis are therefore based on numerical optimization (and so were checked extensively to ensure they are representative of the whole parameter space).

## 2.3 Technological maturity and the no-learning limit

### 2.3.1 Markowitz portfolios

Consider briefly the topic of Markowitz portfolio analysis for standard financial assets (Markowitz 1952). Let  $\mathbf{r} = (r_1, \dots, r_n)^T$  be a vector of stochastic returns (possibly correlated) and  $\mathbf{w} = (w_1, \dots, w_n)^T$  be a vector of portfolio weights. The portfolio return distribution is  $V(\mathbf{w}) = \mathbf{w}^T \mathbf{r}$ , on which a mean-variance optimization is carried out, with  $\mathbf{w}$  as control variable. The classic form of the problem is

$$\begin{aligned} \underset{\mathbf{w}}{\text{maximize:}} \quad & f(\mathbf{w}) = \mathbb{E}[V(\mathbf{w})] - \lambda \text{Var}(V(\mathbf{w})) \\ \text{subject to:} \quad & \sum_{j=1, \dots, n} w_j = 1. \end{aligned} \quad (9)$$

Since this is a mean-variance optimization it looks very similar to our technology portfolio problem (7). There are several differences though; three are superficial but one is fundamental.

First, in the Markowitz case the decision maker seeks high expected portfolio return and low variance, while in the technology case they seek low expected portfolio cost

and low variance<sup>13</sup> — hence the sign difference of the variance terms in (7) and (9). Second, short-selling is in general allowed, so portfolio weights  $w_j$  are not restricted to being non-negative. Third, returns are generally assumed to be correlated, and a lot of attention is paid to understanding these correlations. Finally though, the fundamental difference between the two problems is that in the Markowitz case asset returns are purely stochastic, so portfolio weights do not affect asset performance, while in the experience curve model the stochastic costs depend explicitly on production, so portfolio weights *do* affect technology performances. The more one invests in a given technology the better it gets, on average; there is nonlinear feedback in the technology portfolio model but not in the Markowitz model.

### 2.3.2 Comparing financial and technology portfolios

To better understand the differences between the two portfolio types, we make a more accurate comparison by using a restricted version of the Markowitz model: the no short-selling, enforced budget, uncorrelated, two-asset model. This is a direct equivalent of our technology portfolio problem in a standard financial setting. It eliminates the second and third superficial differences listed above, making it easier to observe feedback effects.

Suppose there are two assets,  $A$  and  $B$ , with uncorrelated normal returns  $r^A \sim \mathcal{N}(\mu^A, (s^A)^2)$  and  $r^B \sim \mathcal{N}(\mu^B, (s^B)^2)$ . Then let  $q^A$  and  $q^B = (1 - q^A)$  be the proportion of wealth invested in  $A$  and  $B$  respectively, with  $q^A, q^B \in [0, 1]$  (the no short-selling condition). The portfolio return distribution is then  $V(q^A) = \sum_{i=A,B} r^i q^i$ , and the objective function to be maximized is

$$f(q^A) = \mathbb{E}[V(q^A)] - \lambda \text{Var}(V(q^A)) \quad (10)$$

$$= \sum_{i=A,B} \mu^i q^i - \lambda (s^i q^i)^2. \quad (11)$$

Note that this is quadratic in portfolio weights  $q^i$ . Then returning to the technology portfolio problem and considering the role of demand  $K$  and initial cumulative productions  $z_0^A$  and  $z_0^B$  in the objective function, a simple calculation reveals the connection between the financial and technology models. Observe that the technologies objective function, Eq. (8), may be written

$$f(q^A) = \sum_{i=A,B} \frac{c_0^i q^i}{(1 + \frac{q^i}{z_0^i})^{\alpha^i}} e^{(\sigma^i)^2/2} + \lambda \left( \frac{c_0^i q^i}{(1 + \frac{q^i}{z_0^i})^{\alpha^i}} \right)^2 e^{(\sigma^i)^2} (e^{(\sigma^i)^2} - 1). \quad (12)$$

When  $q^i/z_0^i$  is small we can approximate this in a simpler form. If the maximum future production of technology  $i$  is much less than its current cumulative production ( $K \ll z_0^i$ ), then  $q^i/z_0^i \ll 1$ , and the binomial series representation

$$(1 + \frac{q^i}{z_0^i})^{-\alpha^i} = 1 - \alpha^i \frac{q^i}{z_0^i} + \alpha^i(\alpha^i + 1) \left( \frac{q^i}{z_0^i} \right)^2 + \dots \quad (13)$$

may be used. Thus if  $K \ll z_0^i$  for both technologies then to zeroth order the objective function may be approximated as

$$f(q^A) \approx \sum_{i=A,B} c_0^i e^{(\sigma^i)^2/2} q^i + \lambda (c_0^i)^2 e^{(\sigma^i)^2} (e^{(\sigma^i)^2} - 1) (q^i)^2, \quad (14)$$

---

<sup>13</sup>This is also the case in the optimal liquidation problem, see e.g. Almgren & Chriss (2001).

which no longer includes the experience exponents  $\alpha^i$ . Appendix A shows details of the expansion, plus higher order terms. Apart from the sign difference of the variance component, this has the same form as the Markowitz model (Eq. (11)), i.e. it is quadratic in production<sup>14</sup>. In this limit learning plays no part, and there is no feedback process by which production affects future costs (since this is represented by the higher order terms). Hence a Markowitz-like portfolio problem is the limiting case of the Wright’s law portfolio problem as learning effects tend to zero.

Eq. (13) shows that a low-learning regime can exist in two ways for a given technology: first, if its learning rate is intrinsically small, and second, if its initial cumulative production is very large compared to the total demand. The latter condition is problem-specific, since it depends on  $K$ , not just on the technology itself. All else being equal, as  $K \rightarrow 0$  technologies behave increasingly like standard financial assets, as noise increasingly dominates learning effects. Furthermore, very mature technologies automatically behave like standard financial assets in the model (since the incremental gains due to learning decrease with maturity by definition in Wright’s law). Note that this analysis relies on the assumption that model parameters are static, e.g. experience exponents are constant and do not depend on the size of  $K$ . This assumption would require justification in any practical application, and indeed is closely related to the question of whether a single- or multi-period framework is more appropriate (the latter could allow for a more fine-grained approach to modelling technological maturity, for example). Nevertheless the simple analytical connection between the two portfolio systems shown here is of note, as it reveals an interesting perspective on technological maturity in a portfolio setting.

Within any given problem then, each technology lies somewhere on a spectrum between more *technology-like* and more *asset-like*, depending on the entire set of parameters. We use “asset-like” simply to mean that learning effects are negligible relative to noise, as with standard financial assets.

Finally, consider how the learning and non-learning portfolio problems differ analytically at lowest orders. As shown in Appendix A, the approximation to  $f$  including the lowest order “learning” terms (i.e. terms in  $-\alpha^i \frac{q^i}{z_0^i}$ ) is

$$f(q^A) \approx \sum_{i=A,B} c_0^i e^{(\sigma^i)^2/2} \left(1 - \alpha^i \frac{q^i}{z_0^i}\right) q^i + \lambda (c_0^i)^2 e^{(\sigma^i)^2} \left(e^{(\sigma^i)^2} - 1\right) \left(1 - 2\alpha^i \frac{q^i}{z_0^i}\right) (q^i)^2. \quad (15)$$

Thus the most straightforward effect of learning is to reduce both the expectation and variance components linearly in  $\alpha^i$ , so that the technology with higher  $\alpha^i$  will perform relatively better in the optimization. In addition though, observe that while the zeroth-order approximation to  $f$  (Eq. (14)) is quadratic in  $q^A$ , and hence always has just one single minimum, the first-order approximation to  $f$  is cubic in  $q^A$ , and may therefore have two local minima inside the optimization range (depending on parameters). The introduction of learning therefore corresponds to the introduction of multiple local optima of the objective function.

---

<sup>14</sup>Note that the  $\sigma^i$  terms remain in the expectation component here due to the particular noise model used (Eq. (4)). They are fixed and independent of  $q^i$  (and indeed could be avoided with a different choice of noise), so do not affect the argument.

### 3 Optimization results

The goal here is to understand how the optimal allocation of production between the two competing technologies depends on the technology-specific learning parameters (experience exponent  $\alpha$  and volatility  $\sigma$ ) and the initial conditions (cost competitiveness  $c_0$  and cumulative production  $z_0$ ) under varying levels of risk aversion  $\lambda$ , for fixed demand  $K$ . To do this we first hold all model parameters constant, then vary technology  $B$  experience exponent  $\alpha^B$  and risk aversion  $\lambda$ . This generates a grid of tuples  $(\alpha^B, \lambda)$ . At each point of this grid the optimization (7) is performed, and the resulting collection of optima is plotted, giving the surface of optimal production of technology  $A$  as a share of total production,  $q_*^A/K$ . The whole process may then be repeated for each of the other technology parameters  $\sigma^B$ ,  $c_0^B$  and  $z_0^B$ .

#### 3.1 Effects of experience exponents $\alpha$

We set the initial conditions and parameter values to those shown in Table 1. Almost identical technologies are used here as this allows us to understand the effects of varying different parameters most effectively. (Asymmetrical technologies are considered in Section 5.) Note that the total demand is twice the initial cumulative production of each technology. Hence the technologies are relatively immature, in the sense that there is plenty of potential left for learning to take place relative to how much has occurred in the past. As shown above, this is necessary since if both technologies are sufficiently mature a nearly-Markowitz scenario emerges.

Fig. 1 shows the surface of optimal technology  $A$  production share,  $q_*^A/K$ , over a grid of  $\alpha^B$  and  $\lambda$  values. This shows how risk aversion and relative experience exponents affect the composition of the optimal portfolio.

Symbol	Description	Tech $A$	Tech $B$
$z_0$	Technology maturity	1	1
$c_0$	Initial cost	2	2
$\alpha$	Experience exponent	0.5	[0-1]
$\sigma$	Technology volatility	1.0	1.1
$K$	Demand	2	
$\lambda$	Risk aversion	[0-1]	

Table 1: Parameter values for the case of two almost identical technologies, as used in Fig. 1




Figure 1: Surface of the optimal production in technology  $A$  as a share of total production. This shows how the optimal portfolio varies with risk aversion  $\lambda$  and technology  $B$  experience exponent  $\alpha^B$  (with  $\alpha^A$  fixed at 0.5). Red areas correspond to higher production of technology  $A$  being optimal, and blue areas to higher production of technology  $B$  being optimal. Low risk aversion leads to more specialized portfolios and greater parameter sensitivity (represented by the surface discontinuity), while high risk aversion leads to greater diversification and lower parameter sensitivity. Parameter values are shown in Table 1.

When risk aversion is low the optimal strategy is to concentrate production entirely in either  $A$  or  $B$  (the dark red and blue plateau regions), depending on relative experience exponents. When risk aversion is high portfolios are diversified over both technologies. This is consistent with a general understanding of both deterministic experience curves (in which specialization is always optimal) and standard portfolio theory (in which diversification reduces portfolio risk). However, the nature of the transitions between these regimes depends on model parameters and is of great interest. For low to moderate risk aversion there is a discontinuity in the surface, indicating a region of extreme sensitivity to model parameters. In this region an incremental change in either experience exponent or risk aversion can lead to a large change in the optimal portfolio. In contrast, for high risk aversion the surface is smooth, so the optimal portfolio is robust to small changes in parameters. As we shall see (in Section 3.3.2), this is caused by the existence of multiple local minima of the objective function in the low risk aversion regime, and a single global minimum in the high risk aversion regime.

On the  $\lambda = 0$  boundary, variance terms do not feature in the optimization so production is concentrated in the technology with the best expected outcome. As risk aversion increases, up to around 0.2, the asymmetry in noise variance becomes apparent and the threshold for switching from 100%  $A$  to 100%  $B$  gradually shifts to larger  $\alpha^B$  values. The preference for the higher experience exponent technology ( $B$  in this region) is traded off




Figure 2: The Markowitz portfolio analogue of the technology portfolio surface shown in Fig. 1. This is the surface of optimal investment share in asset  $A$  for varying values of risk aversion and asset  $B$  expected return. Portfolios are more diversified for high risk aversion and more specialized for low risk aversion as before, and there still exist regions of full specialization, in which one technology sufficiently outperforms the other. However, in contrast to the case of technologies the surface is continuous  $\forall \lambda > 0$ , due to the convexity of the objective function.

against a preference for the less noisy technology ( $A$ ), since the optimization penalizes higher noise variance. As risk aversion increases further portfolios become increasingly balanced. The surface discontinuity becomes less pronounced as the two local optima on either side of it approach a common value. Eventually the discontinuity disappears, and a single stable global optimum exists thereafter. (Only  $\lambda = 0$  is a strict boundary in the model, and the surface extends in the other directions beyond the bounds shown.)

Therefore some combinations of technologies and risk preferences are more robust than others: in some regions the solution is not particularly sensitive to changes in the underlying parameters, while for others it is extremely sensitive. In the unstable regions, a parameter estimation error could lead to a mix of technologies being chosen that is very far from the true optimal mix.

### 3.2 Comparison with Markowitz portfolios

To illustrate how nonlinearities in the technology portfolio affect the optimization results relative to the financial assets case, we plot the corresponding surface of optimal portfolio weights for the equivalent Markowitz system, Eq. (10). With model parameters in Eq. (11) set to  $\mu^A = 0.5$ ,  $s^A = 1.0$ ,  $s^B = 1.1$ , Fig. 2 shows the surface of optima over a grid of varying asset  $B$  expected return  $\mu^B$ , and risk aversion  $\lambda$ . The usual patterns are present: portfolios are more diversified for higher risk aversion and more specialized for

lower risk aversion. Full specialization occurs when one asset sufficiently outperforms the other, again giving the dark red and blue plateau regions. However the crucial difference is that now the surface is continuous everywhere except at the single point on the  $\lambda = 0$  boundary where  $\mu^A = \mu^B$ . There are no positive values of risk aversion at which portfolios transition instantaneously from one state to another; portfolios vary continuously with both risk aversion and model parameters. This is because the Markowitz problem is convex. Without the Wright's law nonlinearity in  $f$  there do not exist multiple local minima for portfolios to instantaneously switch between as parameters vary, and hence no unstable regions of parameter space.

### 3.3 Analysis

Next we present some analytical observations which help in understanding the character of the problem and the shape of the surface in Fig. 1.

#### 3.3.1 Corner and interior solutions

Since the optimization domain is bounded ( $q^A \in [0, K]$ ), solutions are either corner solutions or interior solutions. Corner solutions ( $q_*^A = 0$  or  $K$ ) satisfy  $f'(q_*^A) \neq 0$  almost everywhere in parameter space, while interior solutions ( $q_*^A \in (0, K)$ ) always satisfy  $f'(q_*^A) = 0$ . Corner solutions form both the dark red horizontal plateau with  $q_*^A = K$  on the left of Fig. 1 and the dark blue horizontal floor section with  $q_*^A = 0$  at the front of the plot (plus the equivalent areas on Fig. 2). All other points of the surface are interior solutions, at which optimal portfolios are diversified.

#### 3.3.2 Local and global minima of the objective function

The nonlinearity in the model generates interesting behaviour because in some regions of parameter space the objective function has multiple local optima. Fig. 3 shows how the objective function varies along one particular line in parameter space: risk aversion is fixed at  $\lambda = 0.25$  and technology  $B$  experience exponent is varied (so this corresponds to a section through Fig. 1). The objective function is plotted for three different values of  $\alpha^B$ , showing how distinct local minima emerge and disappear. As  $\alpha^B$  varies the global minimum switches from one local minimum to another, and very different portfolios of approximately equal objective value exist simultaneously. When the surface discontinuity in Fig. 1 is crossed the global minimum switches from one local minimum to the other. This means that a parameter estimation error could lead to a portfolio significantly different to the correct optimal portfolio being chosen. Fig. 4 plots the locations of the different optima against  $\alpha^B$ . This shows how, if the measured value of  $\alpha^B$  is, for example,  $0.7 \pm 0.02$ , then the optimal production share is roughly a 20:80 split, but either technology could be the dominant one, depending on what the true value really is.

Finally, since Fig. 4 is just the  $\lambda = 0.25$  section through Fig. 1, it is apparent that if all optima were plotted on Fig. 1, not just the global minima, the surface would double back under itself in a fold, smoothly connecting the upper and lower edges of the discontinuity. This type of geometry is well-known from the cusp catastrophe bifurcation (see e.g. Zeeman (1976), Poston & Stewart (2014)). Although our setting is different, since parameters here are not dynamic, the similarity is worth noting; both involve plotting the zeros of an underlying nonlinear system, resulting in a multivalued surface representing






Figure 3: The objective function for three different technology  $B$  experience exponents (emphasized by writing  $\alpha^B$  as an argument of  $f$  here). Minima are shown in red, risk aversion is fixed at  $\lambda = 0.25$  and all other parameters are as before. For smaller  $\alpha^B$  there is a single interior local minimum with production concentrated mainly in  $A$ . As  $\alpha^B$  increases a second local minimum appears, which then becomes the global minimum, and production switches to being mainly concentrated in  $B$ . This is what happens as the surface discontinuity in Fig. 1 is crossed — highly differentiated portfolios of approximately equal objective value exist simultaneously.




Figure 4: Locations of the optima of the objective function for varying  $\alpha^B$ , corresponding to Fig. 3. This is the  $\lambda = 0.25$  section through Fig. 1. Distinct local minima emerge and disappear as  $\alpha^B$  varies. At the critical value  $\alpha^B_{switch} \approx 0.71$  the global optimum switches instantaneously between the two minima.




Figure 5: Surfaces of optimal technology  $A$  production share, analogous to Fig. 1, but for varying risk aversion  $\lambda$  and technology  $B$  parameters  $\sigma^B, c_0^B, z_0^B$ .

alternative stable states. Appendix B describes some basic properties of the system that can be derived analytically.

### 3.3.3 Effects of other technology parameters $c_0, z_0$ and $\sigma$

Of the four technology parameter pairs  $\{c_0^i, z_0^i, \alpha^i, \sigma^i\}_{i=A,B}$ , so far we have only studied one: the experience exponents  $\alpha^i$ . To do this we held all parameters constant, including  $\alpha^A$ , then solved the optimization over a grid of varying  $\alpha^B$  and  $\lambda$  values. To study the other three parameter pairs the same procedure is repeated for each of them in turn. Fig. 5 shows the results, with  $\alpha^B = 0.65$  and all other fixed parameters set to the values in Table 1. The plots appear reversed relative to Fig. 1 because while higher values of  $\alpha^i$  correspond to better performance in the model, the opposite is true for the other parameters (e.g. higher values of  $\sigma^i$  are penalised more). Thus we see that all technology parameters produce the same effect as the experience exponents, each pair having its own distinct regions of stability and instability in parameter space. This makes sense intuitively by considering Fig. 3 — at any point in parameter space the objective function is a curve similar to these, and perturbing *any one* of the underlying parameters will cause a similar smooth change in the curve, and its optima, just as it does for the  $\alpha^i$ . The only remaining parameter in the model is demand  $K$ , which we examine next.

## 3.4 Effect of total demand; demand-driven lock-in

In the example system used so far, total demand  $K$  is twice the initial cumulative production of both technologies ( $K = 2z_0^A = 2z_0^B$ ). The potential for learning is therefore high and the model behaves very differently than the Markowitz case. We now consider how this behaviour changes as demand is varied. Intuitively we would expect that, all else being equal, the larger  $K$  is, the more potential there is for experience to accrue, so the more concentrated the portfolio will be in one technology. We demonstrate that our model produces this behaviour, and show in detail the transition from the small- $K$ , low-learning regime to the large- $K$ , high-learning regime.

In Fig. 6, technology  $A$  experience exponent is still fixed at  $\alpha^A = 0.5$ , and instead of varying  $\alpha^B$  as before we also fix it, at  $\alpha^B = 0.65$ , and vary  $K$ . Risk aversion is fixed at  $\lambda = 0.25$  and other parameters are those shown in Table 1 as before (so technology  $B$  progresses faster than  $A$  but is still slightly noisier). The plot shows how the local and global optima of the objective function vary with  $K$  while all other parameters are held




Figure 6: Plot showing how optima of the objective function vary with total production  $K$ . The minimum of the Markowitz approximation to  $f$  (Eq. 14) is also shown for comparison. As  $K$  goes from approximately 0 to 1 the system transitions from a Markowitz-like, low-learning regime to a technology-like, learning regime. When demand is small (approximately  $K < 4$ ) the more certain, slower progressing technology  $A$  ( $\alpha^A = 0.5$ ) dominates the portfolio, but when demand is high there is enough scope for progress to occur that the noisier, faster progressing technology  $B$  ( $\alpha^B = 0.65$ ) becomes optimal. The transition between these states is instantaneous as the global optimum switches between different local minima of equal objective value. Risk aversion is fixed at  $\lambda = 0.25$  and other parameter values are those shown in Table 1 as before.

constant — it is the analogue of Fig. 4, but with  $K$  as the independent variable. As well as the maximum and minima of  $f$ , the plot also shows the minimum of the Markowitz approximation to  $f$  for comparison.

Again we observe instantaneous switching between portfolio states as demand varies. As  $K$  goes from approximately 0 to 1 the share of technology  $A$  in the optimal portfolio (dashed black line) initially decreases sharply then reverses direction and increases again, due to the increased potential for learning. Here the technologies transition from behaving in a more “asset-like” way to a more “technology-like” way, as more higher order terms in the series expansion of  $f$  start to have an impact (see Eq. (15)). Indeed, for very small  $K$ , the minimum of  $f$  and the minimum of the Markowitz approximation to  $f$  roughly coincide (i.e. the black and yellow dashed lines), and the solution for technologies is almost the same as for financial assets. But as  $K$  increases the two curves diverge due to the increasing impact of the nonlinearities in  $f$ , eventually resulting in the appearance of a second minimum when  $K$  is just over 3.

When  $K$  is between approximately 1 and 4, despite technology  $B$  having a larger experience exponent, demand is still too low for it to make enough progress along its experience curve to outweigh its higher variability, so technology  $A$  dominates. But as  $K$  increases a threshold is crossed ( $K \approx 4$ ), and production suddenly switches to  $B$ . Thus we observe a demand-driven *unlocking* of technological lock-in, in the sense that if only a small amount of future demand is considered then it is optimal to continue investing in the slower progressing, less uncertain technology, but if market size is large enough then it is optimal to switch to the noisier, faster progressing technology now.




Figure 7: The objective function for three different values of noise correlation  $\rho$ . Minima are shown in red, risk aversion is fixed at  $\lambda = 0.25$ , technology  $B$  experience exponent is  $\alpha^B = 0.7$  and all other parameters are as before. The  $\rho = 0$  line corresponds to the  $\alpha^B = 0.7$  line in Fig. 3. Greater correlation leads to more specialized portfolios and vice versa.

### 3.5 Correlated noise

In real conditions it is plausible that shocks  $\eta^i$  impacting different technologies are not independent. This might be due, for example, to the fact that the same innovations affect both technologies, reducing both costs. To investigate how the optimization changes in this case we assume that the correlation between  $\eta^A$  and  $\eta^B$  is  $\rho$ . Then the covariance term in  $\text{Var}(V(q^A))$  is nonzero, so the term

$$2\lambda q^A q^B c_0^A \left( \frac{z_0^A}{z_0^A + q^A} \right)^{\alpha^A} e^{(\sigma^A)^2/2} c_0^B \left( \frac{z_0^B}{z_0^B + q^B} \right)^{\alpha^B} e^{(\sigma^B)^2/2} (e^{\rho\sigma^A\sigma^B} - 1) \quad (16)$$

must be added to the objective function (Eq. (8)). Fig. 7 shows the objective function for three different values of  $\rho$ . Technology  $B$  experience exponent is fixed at  $\alpha^B = 0.7$  and all other parameters are the same as before, allowing for direct comparison with Fig. 3. Evidently, the effect of increasing the correlation between technology costs is to decrease the benefit of diversification; i.e. for fixed risk aversion, increased correlation results in more specialized portfolios, and decreased correlation results in more diversified portfolios. Indeed, in this example anti-correlation ( $\rho = -0.1$ ) causes  $f$  to become convex, and thus have a unique minimum. (This could otherwise be achieved by increasing  $\lambda$ .)

Fig. 1 may also be recreated using this version of  $f$  (including the correlation term) and the same effect is observed. This is standard behaviour in a single-period portfolio setting, and does not impact our main findings, so we continue to consider only uncorrelated noise henceforth.

## 4 The efficient frontier

Technology portfolios can also be viewed in the efficient frontier framework. This technique is well-known in portfolio theory, and involves plotting each portfolio as a point in expected-return/variance space. We first describe the approach for the restricted Markowitz problem introduced earlier, then show how it applies for technologies (though note that we only consider the no risk-free asset case.)

### 4.1 Financial assets

In the Markowitz system defined above (Eq. (10)) there are two assets, with known return distributions. The portfolio weight of asset  $A$ ,  $q^A$ , is the single free control variable. Each  $q^A \in [0, 1]$  describes a unique portfolio and as  $q^A$  varies from 0 to 1 all feasible portfolios are spanned. For a given value of risk aversion only one of these portfolios is optimal. Each portfolio has a return distribution  $V(q^A)$ , the expectation and variance of which may be used to plot a single point in expectation-variance space representing the portfolio. This gives the well-known Markowitz diagram: the  $x$ -axis is the portfolio variance,  $\text{Var}(V(q^A))$ , and the  $y$ -axis is the expected return,  $\mathbb{E}[V(q^A)]$ . The feasible set of portfolios is the curve traced out on these axes as  $q^A$  varies from 0 to 1. Fig. 8 shows the feasible set for two assets with fixed parameters  $\mu^A = 0.5$ ,  $\mu^B = 0.65$ ,  $s^A = 1.0$  and  $s^B = 1.1$  (*cf.* Section 3.2). This horizontal parabola is known as the Markowitz bullet. The colour scheme is the same as before (*cf.* Fig. 2) so dark red corresponds to 100% asset  $A$  and dark blue to 100% asset  $B$ .

From the definition of  $f$  (Eq. (10)) we have  $\mathbb{E}[V] = f(V) + \lambda \text{Var}(V)$ , so on these axes the isolines of  $f$  (i.e. level sets of  $f$ ), for any fixed value of risk aversion, are just the straight lines of gradient  $\lambda$ . The value of  $f$  for each portfolio lying on a given isoline (i.e. the points of intersection with the feasible set) is given by the  $y$ -axis intercept. Then since we want to *maximise*  $f$  in this problem, the optimal portfolio for this  $\lambda$  is the unique point of intersection of the feasible set and the isoline of gradient  $\lambda$  with highest  $y$ -axis intercept. The *efficient frontier* is defined as the set of all portfolios which are optimal for some value of  $\lambda$ . Therefore, since  $\lambda \in [0, \infty)$ , the efficient frontier here is the segment of the feasible set furthest into the upper-left-most quadrant of the diagram. These elements are all shown on Fig. 8.

### 4.2 Technologies

In contrast to the Markowitz model, in the technologies model we want to minimize both the variance *and* the expected cost, so the sign of the variance part of the objective function is reversed. The isolines of  $f$  are therefore now the *downward*-sloping straight lines, of gradient  $-\lambda \in (-\infty, 0]$ . And since lower  $f$  is now better, an optimal portfolio is a point of intersection of the feasible set with the isoline of *lowest*  $y$ -axis intercept. The efficient frontier therefore consists of the part of the feasible set furthest into the lower-left-most quadrant of the diagram. We demonstrate these differences between the Markowitz and technology models by first plotting the expectation-variance diagram for two technologies in a low-learning regime, and then for the same two technologies in a high-learning regime. These are shown in Figs. 9 and 10, and are analogous to Fig. 8.

In Fig. 9, technology  $B$  has experience exponent  $\alpha^B = 0.65$  and other parameters are those shown in Table 1 as before, except for demand, which is set to  $K = 0.1$ . By




Figure 8: The feasible set of portfolios for the Markowitz system Eq. (10), with  $\mu^A = 0.5$ ,  $\mu^B = 0.65$ ,  $s^A = 1.0$  and  $s^B = 1.1$ . This is the path of portfolios traced out as the proportion of asset  $A$  in the portfolio varies from 0% (dark blue,  $q^A = 0$ ) to 100% (dark red,  $q^A = 1$ ). To demonstrate how risk aversion and optimality are related geometrically, an isoline of  $f$  for risk aversion  $\lambda = 0.25$  is plotted. The black dot at the point of tangency with the feasible set is the unique optimal portfolio for this  $\lambda$ . The two other black dots represent the two full specialization portfolios.




Figure 9: The feasible set of portfolios for two technologies in a low-learning regime. This is the path of portfolios traced out as the proportion of technology  $A$  production in the portfolio varies from 0% (dark blue,  $q^A = 0$ ) to 100% (dark red,  $q^A = K$ ). The technologies here have  $\alpha^A = 0.5$ ,  $\alpha^B = 0.65$ , and other parameters are those shown in Table 1, except demand, which is set to  $K = 0.1$ . This severely limits the potential for learning, so the problem is nearly-Markowitz and hence the feasible set is almost parabolic. Isolines of  $f$  now slope downward and the efficient frontier is the lower-left-most portion of the feasible set. An isoline corresponding to risk aversion  $\lambda = 0.25$  is plotted. The black dot at the point of tangency with the feasible set is the unique optimal portfolio for this  $\lambda$ .

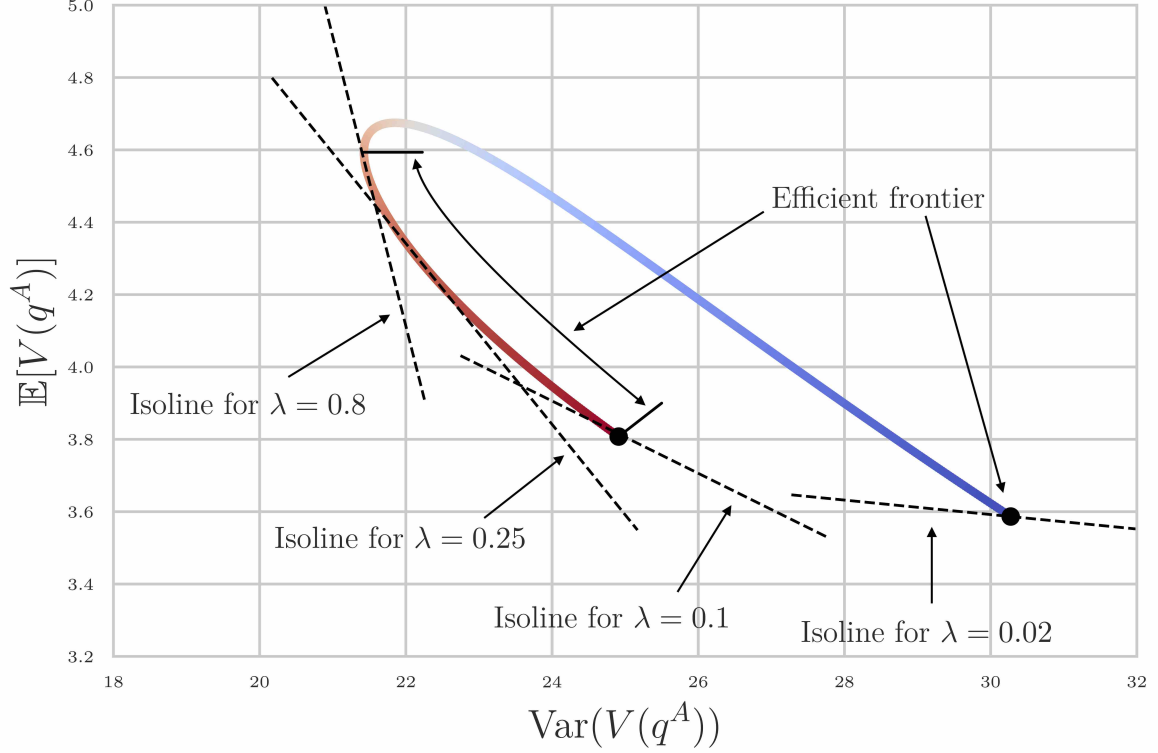


Figure 10: The feasible set of portfolios for technologies analogous to Fig. 9, but with demand now set at the higher value  $K = 2$ . The technologies are no longer in a low-learning regime, and nonlinear feedback causes the feasible set to be stretched and tilted. Isolines for four values of risk aversion are plotted, showing how the transformed geometry of the feasible set causes the efficient frontier to now consist of two disconnected components. There is a critical value of risk aversion ( $0.02 < \lambda_{switch} < 0.1$ ) at which two optimal portfolios exist simultaneously (100% tech  $A$  and 100% tech  $B$ ). This is how instantaneous optimum switching is manifested in expectation-variance space.



restricting production in this way, learning effects are very small, so we are in a nearly-Markowitz regime, and hence the feasible set is almost parabolic. The figure makes clear how the efficient frontier and isolines differ in the technologies problem, as compared to the Markowitz problem (Fig. 8). In Fig. 10 all parameters are identical to Fig. 9, except demand, which is now returned to the value in Table 1,  $K = 2$ , so that the technologies are no longer in a low-learning regime.

As these plots show, there are two significant differences between how technologies and financial assets appear in this framework. First, for technologies the feasible set is tilted and stretched. This is because the objective function is highly nonlinear in portfolio weights, not just quadratic. Second, as a direct consequence of this, the efficient frontier may now be split into two disconnected components. This is the case in Fig. 10, where the efficient frontier consists of both the long red segment on the left (mainly technology  $A$ ) and the isolated end-point on the right (100% technology  $B$ ). This splitting of the efficient frontier is how instantaneous optimum switching is manifested in expectation-variance space: as risk aversion goes from  $\lambda = 0$  to  $\infty$  the optimal portfolio traverses the efficient frontier from one end to the other, jumping from one component to the other at the critical value  $\lambda = \lambda_{switch}$ . At the point of the discontinuity  $f$  has two distinct minima of equal value, and there are two optimal portfolios, both of which lie on the same isoline, of gradient  $-\lambda_{switch}$ . In this case these are the 100%  $A$  and 100%  $B$  portfolios.

Note that since  $\alpha^B = 0.65$  in Fig. 10, the efficient frontier shown corresponds to the  $\alpha^B = 0.65$  section through the surface in Fig. 1. Hence the change in the optimal portfolio as risk aversion varies can be traced out equivalently on both diagrams. The closed-form expression for  $\lambda_{switch}$  is given in Appendix B.2.

Viewing the problem in the expectation-variance framework shows how optimal technology portfolios of equal value can coexist simultaneously (unlike in financial portfolios). Similar value portfolios may have either large expectation and small variance, or vice versa, or some combination in between. And since the feasible set is no longer parabolic (due to Wright’s law nonlinearities), there may be many very different portfolios lying near the optimal isoline. For example, in Fig. 10 all portfolios with around 60-90% technology  $A$  (red) lie very near to the optimal  $\lambda = 0.25$  isoline, because the feasible set has very low curvature here.

This is suggestive of the behaviour we would expect to see in a multi-technology model. The increasing returns dynamic allows many different ways of generating portfolios of similar value, using different combinations of the various technologies’ expectations and variances. This would result in a highly non-convex optimization problem with many local minima.

### 4.3 Effect of demand on the efficient frontier

Fig. 11 shows in more detail the effect of total demand on the feasible set and optimal portfolios. The technologies are fixed, and are the same as in Figs. 9 and 10.  $K$  is the only parameter which varies. Although the scales on the axes differ in each plot, the ratio between them is constant. Indeed, the dashed lines all have gradient  $\lambda = -0.25$ ; they are the isolines corresponding to the optimal portfolio for  $\lambda = 0.25$  in each case.

When  $K$  is tiny there is very little potential for learning so the technologies behave in an “asset-like” way, and the feasible set is almost a parabolic Markowitz bullet. Here technology  $A$  (red) dominates the optimal portfolio for all values of risk aversion. As  $K$  increases the potential for learning increases, so the nonlinearities in  $f$  start to have an

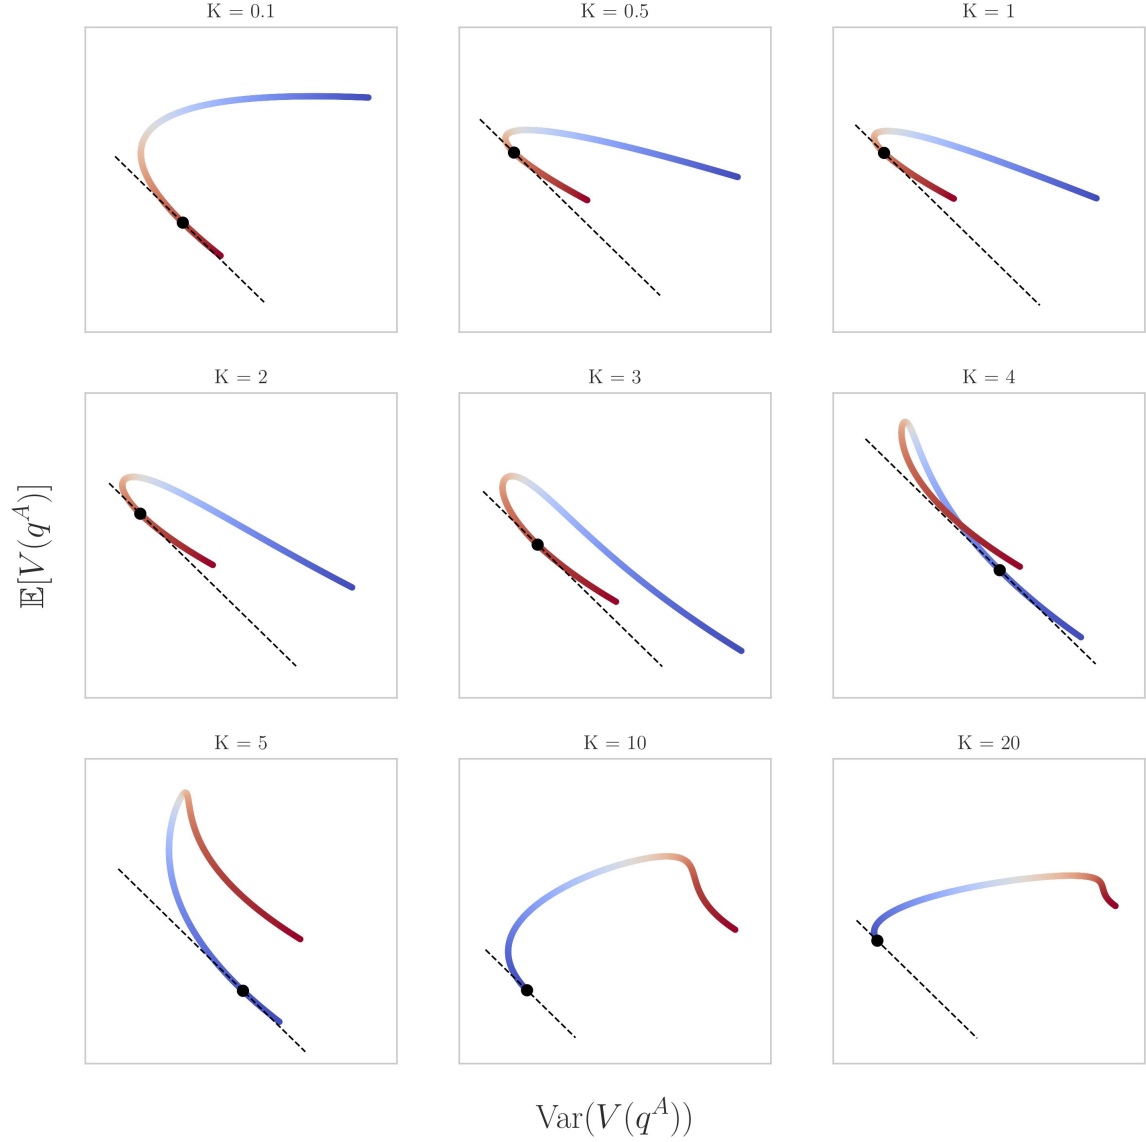


Figure 11: Feasible portfolios for  $\alpha^A = 0.5$ ,  $\alpha^B = 0.65$  and varying  $K$ , with other parameters fixed as before (Table 1). Again, dark blue corresponds to 0% technology  $A$  ( $q^A = 0$ ) and dark red to 100% ( $q^A = K$ ). The dashed lines are  $\lambda = 0.25$  isolines of  $f$ , and the black dots are the optimal portfolio for this risk aversion. The axes scales have been omitted for clarity (they are different for each plot). These plots show how the problem transitions from a Markowitz-like, low-learning regime when  $K$  is small, to a highly nonlinear high-learning regime when  $K$  is large.

impact and the feasible set starts to become distorted. At around  $K = 1$  the blue arm drops below the red arm, splitting the efficient frontier in two and indicating the presence of two equal value portfolios for the first time. Here, for  $\lambda \leq \lambda_{switch}$  100% technology  $B$  (blue) is optimal, but for  $\lambda \geq \lambda_{switch}$  technology  $A$  (red) is dominant. At around  $K = 4$  the two arms cross and large portions of the feasible set are very close to the  $\lambda = 0.25$  isoline. Hence there are many different nearly-optimal portfolios in this case. After this point the two arms of the feasible set cross over completely so that technology  $B$  is dominant for all values of risk aversion. As  $K$  gets very large ( $> 10$ ) the potential for learning is so great that near full specialization in the fastest-learning technology ( $B$ ) is optimal for all levels of risk aversion. Fig. 11 may be related back to Fig. 6 (though note that distance increments *along* the feasible set do not correspond linearly to increments in  $q^A/K$ ).

Clearly this analysis relies on the assumption that noise is independent of total production  $K$ , so that more production, and hence learning, can take place without affecting the size of the shocks.

## 5 Asymmetrical technologies and escaping lock-in

The paper so far has focused on the case of two almost symmetrical technologies, studying the behaviour of the optimal portfolio as one parameter is changed, while all others are held constant. We now consider a more realistic and interesting example in which an established technology is challenged by a newcomer<sup>15</sup>. Suppose the setting is one where a cheap, mature, slow-learning technology  $A$  is challenged by a costly, young, fast-learning technology  $B$ . Table 2 shows the parameter values used here.

Symbol	Description	Tech $A$	Tech $B$
$z_0$	Technology maturity	100	1
$c_0$	Initial cost	1	2
$\alpha$	Experience exponent	0.15	0.2
$\sigma$	Technology volatility	0.1	0.1
$K$	Demand	[0-100]	
$\lambda$	Risk aversion	[0-1.2]	

Table 2: Parameter values for the case where a young expensive technology competes with an old cheap technology.

<sup>15</sup>The limiting case of this is when one technology is a ‘safe’ technology, with constant cost ( $\alpha^A = \sigma^A = 0$ ). Then, in the nearly-Markowitz regime ( $K \ll z_0^B$ ), the objective function reduces to

$$f \approx c_0^A (K - q^B) + c_0^B e^{(\sigma^B)^2/2} q^B + \lambda (c_0^B)^2 e^{(\sigma^B)^2} (e^{(\sigma^B)^2} - 1) (q^B)^2 \quad (17)$$

(cf. Eq. (14)). This is a convex parabola with minimum

$$q_*^B = \frac{c_0^A - c_0^B e^{(\sigma^B)^2/2}}{2\lambda (c_0^B)^2 e^{(\sigma^B)^2} (e^{(\sigma^B)^2} - 1)}, \quad (18)$$

and the condition for the optimal solution to be diversified between the safe and the new technology is  $0 < q_*^B < K$ . This is analogous to the Markowitz portfolio problem with a safe asset and a risky asset. Outside of this low-learning regime though, the safe technology simplification does not result in increased analytical tractability, so the numerical solution approach is still applicable.




Figure 12: Optimal share of production of incumbent technology  $A$  (red) when in competition with challenger technology  $B$  (blue), for varying demand  $K$  and risk aversion  $\lambda$ . Parameter values are shown in Table 2. Technology  $A$  has such a strong initial cost advantage that when demand is low (and hence the potential for learning is low), it is optimal to specialize fully in  $A$ , for all values of risk aversion shown. As demand increases, so does the potential for learning, and in order to exploit the faster-learning challenger the global optimum switches to a new local minimum, in which technology  $B$  dominates.

Repeating the demand-driven lock-in analysis of Section 3.4, Fig. 12 shows the optimal portfolio surface over total demand and risk aversion axes. Technology  $A$ 's initial cost advantage is so strong that a demand of at least 20 times the initial cumulative production of technology  $B$  is required to prevent full specialization in  $A$ , even for high risk aversion. We observe the familiar optimum switching as demand increases, demonstrating again how important the role of anticipated future demand is in determining the optimal production mix. In this case, the  $K$ - $\lambda$  parameter space separates into three qualitatively different regimes: *i*) for low  $K$  100% technology  $A$  is optimal, *ii*) for high  $K$  but low risk aversion 100% technology  $B$  is optimal, and *iii*) for high  $K$  and at least moderate risk aversion, the optimal proportion of technology  $A$  is around 0-40%. Results like this could be very useful in applications, where often the key challenge is simply reducing the dimension of the decision space.

For large  $K$  there is a qualitative change in the surface at  $\lambda \approx 0.4$ . Here the global optimum of the objective function moves from the boundary to the interior of the optimization range (the value of  $\lambda$  at which this transition occurs may be found analytically as shown in Appendix B.1). This is similar to the situation in Fig. 3, where as  $\alpha^B$  increases from 0.71 to 0.72 up to, say 0.9 (not shown), the global minimum moves towards the boundary, then sticks on it (and in fact the function minimum over  $\mathbb{R}$  continues to move further outside the optimization range, but this is not a valid solution).

For comparison, Appendix C shows the same demand-risk surface for the case of the two almost identical technologies of Section 3.

To summarise the behaviour of the model, note the direction in which each of the parameters would need to move in order to escape lock-in to an incumbent technology. All things being equal, avoiding lock-in to technology  $A$  would require: decreased  $B$  maturity, decreased  $B$  initial cost, increased  $B$  experience exponent, decreased  $B$  volatility, increased demand  $K$ , or increased risk aversion  $\lambda$ .

## 6 Two-period model

Having analysed the model in the simplest, single-period case, we now consider the extension to two periods. This allows us to introduce discounting and investigate its effect on optimal production timing. This is especially relevant for experience curves since they exemplify the concept of investing effort now to unlock future benefits, so the relative value of present and future benefits is critical. Moving to a multi-period setting allows us to investigate the conditions under which we should plan to invest in a technology now, or in future (or neither), given current knowledge about the present state of technologies and their likely development under various investment scenarios. We briefly discuss some features and limitations of our approach.

The extension we consider is static, in the sense that the model schedules production now, *for all future periods*, in the optimal manner as defined by the objective function. This static analysis relies on current estimates of technology parameters, which are in turn based on empirical data about the technologies in question. But as time progresses we observe realisations of the noise, thereby gaining new information about technology costs and parameters. Therefore performing the same optimization procedure in future may yield different results, and we are faced with the possibility of time-inconsistency. This is a limitation of the static model<sup>16</sup>. In support of the static approach however, note that portfolio adjustment costs could be very high for technologies, which may result in high levels of commitment for future periods anyway.

In this model, technological progress only occurs via the stochastic experience curve mechanism, there is no exogenous progress trend. This is due to our choice of zero mean noise,  $\eta \sim \mathcal{N}(0, \sigma^2)$ , which we use because it is close to the model tested empirically by Lafond et al. (2018). By using normal noise with nonzero mean instead it is possible to model an exogenous progress trend. For  $\eta \sim \mathcal{N}(\mu, \sigma^2)$ , the standard lognormal distribution has expectation  $e^{\mu+\sigma^2/2}$ , so if  $\mu < -\sigma^2/2$  then the expected cost of a technology can in fact decrease between periods under zero production (*cf.* Eq. (4)). This is especially important in the multi-period setting (although it also applies in the single-period case). If a technology is initially very expensive, and  $\mu$  is very negative, then waiting for the technology to improve may indeed be a viable strategy. However, while waiting for this improvement in the expectation, the cost variance may become less favourable (relative to that of the other technology), counteracting any benefit in the objective function.

---

<sup>16</sup>The dynamic mean-variance portfolio problem is difficult even in the case of standard financial returns, see e.g. Gârleanu & Pedersen (2013). In our case the complexities are increased by the nonlinearities and time dependencies in costs. One way to approach the problem would be to discretize the decision space and the random event space, then apply backward induction to the resulting “scenario tree”. See, e.g., Edirisinghe & Patterson (2007) for an application of this method to the standard mean-variance portfolio problem. However, even for rough discretization the resulting tree would grow very quickly. Hence, we do not pursue this analysis here, but leave it for future research.

Generally, increasing experience in one technology comes at the expense of the other, and if it is optimal to “delay production” in one technology then this is due only to the complex interplay of all model parameters in the objective function, not because of some simple background improvement in a technology. Here though we continue to use zero mean noise.

For simplicity, as in the one-period case, we use technology production, not investment, as a proxy for experience. In the single period case the distinction is not significant, but when considering multiple periods it is, due to capital depreciation. Thus our multi-period framework does *not* model investment strategies, only production strategies. A proper treatment of investment would need to model the relationship between production and investment, which would require at least one extra parameter to represent depreciation. We prefer not to complicate the model further, and instead just consider production, while highlighting this discrepancy.

## 6.1 Optimization

We use the same first-difference Wright’s law model as before (Eq. (2)), with subscripts now denoting distinct periods, and separate noise shocks impacting each period:

$$\log(c_t^i) - \log(c_{t-1}^i) = -\alpha^i [\log(z_t^i) - \log(z_{t-1}^i)] + \eta_t^i, \quad i = A, B, \quad t = 1, 2. \quad (19)$$

Production of each technology accumulates in the obvious way:  $z_t^i = z_0^i + \sum_{k=1}^t q_k^i$ . We retain the assumption that each technology has its own specific noise distribution,  $\eta^i \sim \mathcal{N}(0, (\sigma^i)^2)$ , and now assume that in each period there is a new draw from this distribution,  $\eta_t^i$ . For simplicity we assume shocks are uncorrelated over both time and technologies (though clearly correlated noise has the potential to play an important role in a multi-period setting). Average unit costs in the second period are then given by:

$$c_2^i = c_1^i \left( \frac{z_1}{z_2} \right)^{\alpha^i} e^{\eta_2^i} = c_0^i \left( \frac{z_0}{z_2} \right)^{\alpha^i} e^{\eta_1^i + \eta_2^i}. \quad (20)$$

This makes it clear how successive shocks impact cost in the two-period setting: from period to period costs fall according to Wright’s law, but the periodic shocks accumulate, potentially driving the cost far from the deterministic experience curve trend. In contrast to standard experience curve implementations, this has the advantage that costs are not guaranteed to fall in the long run due to experience effects, and instead admits the possibility that uncertain exogenous events may dominate.

While it is possible to allow any of the parameters to vary over the periods in this setup (e.g. experience exponents, noise distributions), here we keep them all fixed, again for simplicity and clarity. Total production is set to  $K$  *per period*, so there are now two production constraints:  $q_t^A + q_t^B = K$  for  $t = 1, 2$ . The above choices are made because the model gets unwieldy very quickly, so it is preferable to use the simplest possible formulation that captures the key features of interest.

We implement exponential discounting with discount rate  $r$  and consider the present discounted cost of the total system:

$$V(q_1^A, q_2^A) = \sum_{t=1,2} \sum_{i=A,B} e^{-r(t-1)} c_t^i q_t^i. \quad (21)$$

In contrast to the one-period total cost (Eq. (6)), this is now a function of production in both periods. We use the same mean-variance objective function as before, details of which are given in Appendix D<sup>17</sup>. The optimization problem is then

$$\underset{q_1^A, q_2^A}{\text{minimize:}} \quad f(q_1^A, q_2^A) = \mathbb{E} [V(q_1^A, q_2^A)] + \lambda \text{Var} (V(q_1^A, q_2^A)) \quad (22)$$

$$\text{subject to: } q_1^A \in [0, K], \quad q_2^A \in [0, K].$$

## 6.2 Results

To show how discount rate  $r$  and risk aversion  $\lambda$  interact to influence optimal production over the two periods we fix all other technology parameters and production constraints, then plot the objective function for a range of  $r$ - $\lambda$  pairs. Since  $f$  is now a function of the two control variables,  $q_1^A$  and  $q_2^A$ , we plot its values as a contour-/heat-map over these axes.

Fig. 13 shows the results for nine pairs of low, medium and high discount rate and risk aversion:  $r = 0.1, 1.0, 3.0$ ,  $\lambda = 0.1, 0.5, 3.0$ .<sup>18</sup> The two technologies used here are the same as in the asymmetrical case of Section 5 (Table 2), and total production  $K$  is set to 30 units per period, as this allows results to be related directly back to the one-period case in Fig. 12. (Production is shown here in units instead of percentages for easier comparison.) Darker blue corresponds to lower values of  $f$  and darker red to higher values, though the scale is different for each plot. The global minimum, again obtained by brute force optimization, is shown as a black dot (though this sometimes forms part of a valley of nearly identical values).

## 6.3 Discussion

We observe the same kind of phenomena as in the one-period case: highly specialized solutions for low risk aversion (due to deterministic learning feedback), greater diversification for high risk aversion (due to the counteracting effects of variance terms in  $f$ ), multiple optima of the objective function, and discontinuities in the global optimum as parameters vary.

First note that as the discount rate increases the two-period problem reduces to the one-period problem. Plots in the  $r = 3.0$  column are nearly uniform in the  $q_2^A$  direction, since second period costs are discounted so heavily that the choice of second period production makes little difference to  $f$ . The global optima here correspond to the global optima shown at the relevant points on Fig. 12.

To observe the presence of multiple optima and instantaneous switching of the global minimum, consider the two left-most plots of the top and middle rows of Fig. 13. In the top row ( $\lambda = 0.1$ ), the two plots are formed of concentric oval isolines, with a maximum of  $f$  at their centres. All four corners are local minima, reflecting the fact that strong feedback in the deterministic experience curve setting discourages any diversification at all.

<sup>17</sup>As mentioned earlier, the choice of objective function in a multi-period setting is more subtle and important than for a single period; we leave this consideration for future research.

<sup>18</sup>These values are picked for illustrative purposes only, in order to display various features clearly. However, to justify these discount rates, it is useful to think about the length of one period being about 30-40 years. This could be motivated by long horizons in the energy sector; for instance, thermal power plants are designed for a life of 30 to 40 years.

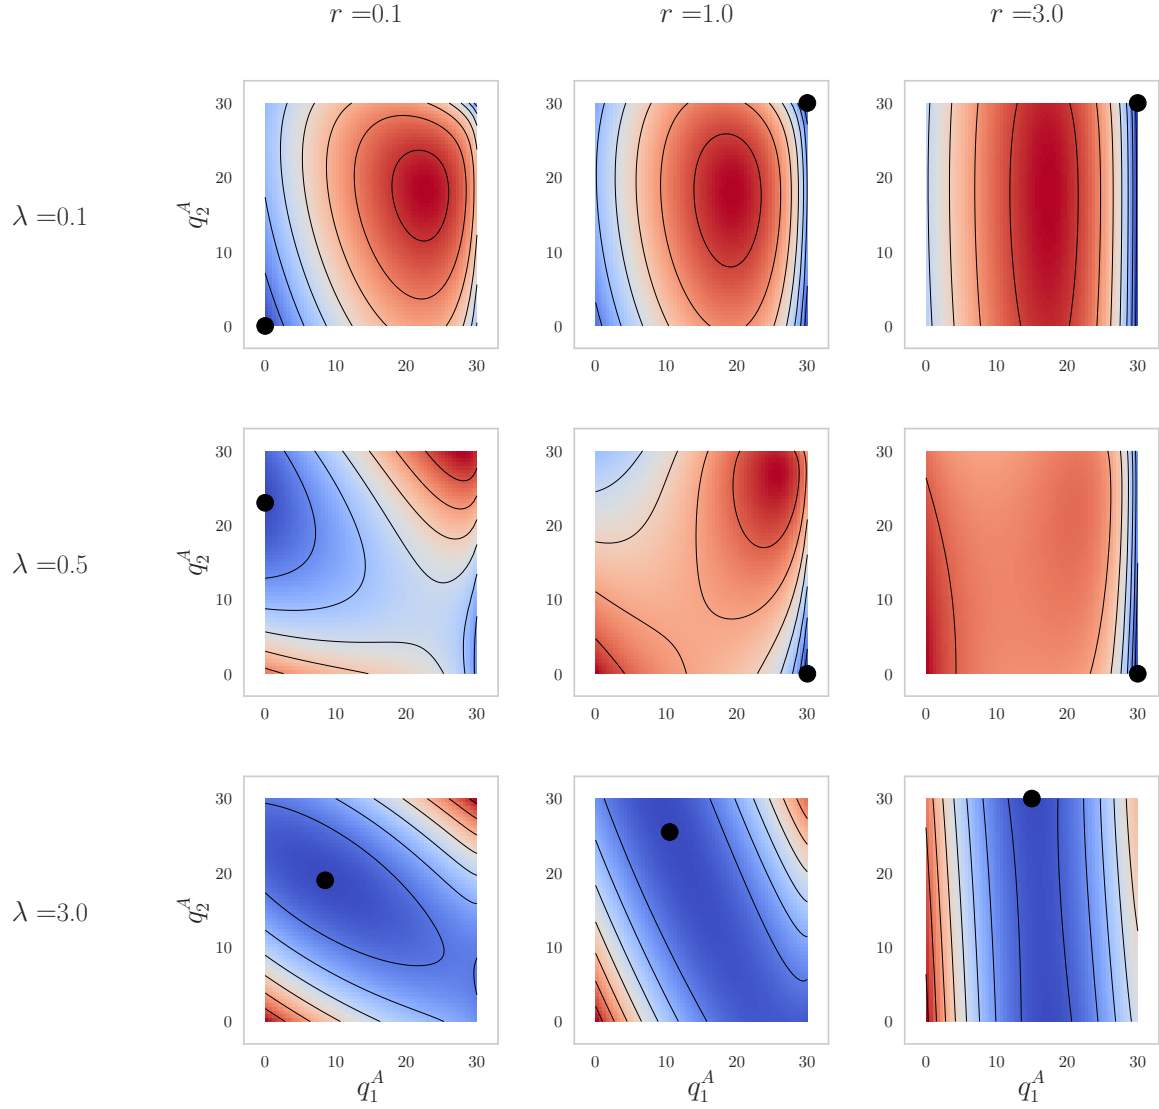


Figure 13: Heat-maps showing the value of the two-period objective function for a range of risk aversion and discount values, as a function of first and second period production in technology  $A$  ( $q_1^A$ ,  $q_2^A$ ). Blue corresponds to lower values of  $f$ , red to higher values, and the global optimum is shown as a black dot. The technologies used here are the asymmetrical technologies of Section 5:  $A$  is a cheaper, mature, incumbent technology and  $B$  is a young, fast-learning challenger.



As  $r$  goes from 0.1 to 1.0, the global optimum switches from (0, 0) to (30, 30) (similarly to Fig. 3). This is because when second period costs are only weakly discounted, technology  $B$ 's higher experience exponent has the potential to generate sufficiently low expected second period cost that it is worth concentrating production in  $B$  in the first period, then reaping the rewards in the second. But when second period costs are moderately discounted, this second period 'reward' has lower overall value in  $f$ , so it is not worth producing any technology  $B$  at all in either period. (This is just standard deterministic experience curve behaviour with discounting.)

In the middle row ( $\lambda = 0.5$ ) we again observe optimum switching behaviour, but the situation is more nuanced.  $f$  is now a saddle, with opposite corners forming distinct minima and maxima as the different expectation and variance components of  $f$  contribute and interact in complex ways. In comparison with the top row though, the central areas of these plots are generally less unattractive, since risk aversion allows more balanced portfolios. As  $r$  varies from 0.1 to 1.0, the global optimum switches from around (0, 23) to (30, 0). When discounting is very weak the optimal strategy is again to concentrate production in technology  $B$  initially in order to bring the cost down, but then to diversify slightly in period two in order to lower the variance. When discounting is moderate though, the expected cost reduction in technology  $B$  is not valuable enough to override the initial cost advantage of  $A$ , so first period production is concentrated in  $A$ . Then however, production switches to being entirely concentrated in technology  $B$  in the second period. This is likely due to the fact that the covariance term in  $f$  (see Eq. (52)) contains the product  $q_1^A q_2^A$ , so there is a benefit to setting one of these terms equal to zero, and in this particular case the effect is large enough to produce the bang-bang solution shown. However, since all 10 components of  $f$  (see Appendix D) are being traded off against each other in this moderate risk aversion/ moderate discounting regime, the relative size of any specific effect is not clear (without further analysis).

Finally, in the bottom row ( $\lambda = 3.0$ ), risk aversion is strong and hence diversification is optimal. Whole regions of the solution space generate identical or similar values of  $f$ . This is because, as in the one-period case, expected progress is being traded off against the certainty in this progress, and the same objective function value may be achieved by many different combinations of first and second period production.

As in the one-period problem, due to the functional form of the objective function, the instantaneous optimum-switching behaviour demonstrated here with the  $r$  and  $\lambda$  parameters can also occur when any one of the underlying parameters is varied continuously.

## 6.4 Effects of discounting and risk aversion in scenario comparison

While optimization over the entire solution space is essential for understanding the basic properties of the model, there are situations in which this is either undesirable or impossible, and it is better simply to compare the performance of a restricted set of portfolios. We refer to this as scenario comparison (as this captures better the idea of comparing worlds characterised by different technology mixes). This can be useful either if there are additional system constraints, or if the solution space is very large, for the following reasons.

First, in practise technologies are not perfect substitutes, so although they may be considered substitutes in terms of some primary characteristic (i.e. "production"), other, secondary characteristics may require portfolios to be constrained in additional ways. For




Figure 14: Plots showing how the objective function varies with the discount rate for each of three different technology scenarios, in low, moderate and high risk aversion regimes. The technologies are the same as in Fig. 13. For low and moderate risk aversion there is a critical discount rate separating regimes in which different scenarios are preferred.

example, solutions that vary wildly from period to period may be impractical in many real world situations. Or as another example, consider a portfolio of energy technologies on an electric power grid. While they may be considered substitutes in terms of, say, total annual energy production, many other engineering and physical constraints must be met in order to ensure grid stability<sup>19</sup>. This would place extra restrictions on such technology portfolios; these constraints could of course be modelled explicitly and a more complicated objective function constructed, but this is outside the scope (and spirit) of the work here.

Second, consider extending the model to include more technologies and/or periods. As the number of control variables increases, brute force optimization quickly becomes computationally intractable, and since the problem is non-convex (due to the experience curve feedback), local optimization methods are not guaranteed to find the global optimum. It may therefore be necessary to use heuristic arguments to specify a restricted set of available portfolios, and compare these directly.

Now, since portfolios are fixed in the scenario comparison setting, this allows for an alternative perspective on the effects of discounting and risk aversion. We demonstrate this with an example, using the same two-period model and asymmetrical technologies as earlier in this section.

Suppose that due to extra system constraints, only three scenarios are available: 95% technology *A* in both periods, 50% technology *A* in both periods and 95% technology *B* in both periods. Fig. 14 shows three different risk aversion regimes ( $\lambda = 0.1, 0.5, 3.0$ ), and in each regime the objective function is shown as a function of discount rate, for each scenario. This shows how the discount rate affects the preference ordering of the three available scenarios in each risk aversion regime. These plots are readily related back to Fig. 13 above.

When risk aversion is low the preferred scenario is 95% technology *B* in both periods for mild discounting, but as discounting increases this becomes less advantageous, and there is a critical discount rate above which the preferred scenario becomes 95% technology *A* in both periods. When risk aversion is moderate, 50% *A* in both periods

<sup>19</sup>E.g. if large shares of intermittent renewable technologies are present then large quantities of energy storage or other backup technologies may also be required, depending on daily and seasonal demand patterns.

is preferred for low discounting, all three scenarios are roughly equal when discounting is moderate, and 95%  $A$  in both periods is preferred when discounting is high. When risk aversion is high a 50-50 mix of technologies in both periods is the preferred scenario regardless of discount rate.

The same arguments and explanations given in Section 6.3 apply here. Clearly this is only a stylized model, with limited scenarios, but the key point is that in some circumstances a critical discount rate exists separating regimes in which different scenarios are preferred. While this type of result is well known in deterministic situations, and seems trivial in a low dimensional stochastic model such as this, in a multi-technology, multi-period setting this technique could provide essential insight into preference orderings of different technology scenarios.

## 7 Conclusion

In this paper we considered two technologies following stochastic experience curves and we characterised the optimal investment (or production) strategies. We used an objective function that accounts for total portfolio cost and uncertainty. The optimal investment depends on risk-aversion, initial conditions (relative technology maturity and initial cost), progress characteristics (mean progress rate and uncertainty of future shocks) and market size; and in the multi-period case also the discount rate.

In contrast to classical Markowitz portfolios, in our setting investment lowers marginal cost, creating a larger and larger incentive to continue investing in the same option. But contrary to the deterministic case, in which these self-reinforcing effects lead to complete specialization, we find that accounting for uncertainty and risk aversion promotes diversification even if one option has better intrinsic technological characteristics. Our results therefore show how the choice of specializing or diversifying depends on the underlying parameters and initial conditions. Crucially, we find that the nonlinearity of the problem leads to multiple local optima, so that very different optimal portfolios can exist simultaneously, and the global optimum switches instantaneously between them as parameters change. This means that inside a critical region of parameter space a small change in one of the parameters can lead to a very significant change in the optimal portfolio.

We established an analytical connection between portfolios of technologies and portfolios of financial assets. The Wright's law model of endogenous technological change may be expanded in a series approximation, the leading terms of which are equivalent to those found in the Markowitz model for financial assets. Only the higher order nonlinear terms contribute towards the learning feedback, and hence the Markowitz model may be regarded as the no-learning limit of the technology portfolio model. We also showed that the strength of learning feedback in the model depends on the complete set of model parameters, but in particular on technologies' previous cumulative production and the level of total future demand. As a result, each technology may be viewed as existing somewhere on a spectrum between more asset-like and more technology-like, depending on the specific set of parameters present.

We also considered the two-period case and found that in a scenario comparison setting (i.e. when comparing a limited number of available portfolios, as would be typical in higher-dimensional applications), the discount rate plays a critical role. For a given level of risk aversion, the discount rate determines which of the available portfolios is preferred.

These findings help establish a theoretical basis for understanding how technologies behave in a simple mean-variance framework, and give insight into how multiple optima can arise in the context of uncertain endogenous technological change.

## Acknowledgements

This project was primarily supported by Partners for a New Economy. The project also received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 730427 (COP21 RIPPLES), and from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ ERC grant agreement No. 611272 (GROWTHCOM). The authors gratefully acknowledge all these sources of financial support, and additionally the Institute for New Economic Thinking at the Oxford Martin School for its continuing support. We would also like to thank Cameron Hepburn, Alex Teytelboym, Torsten Heinrich, Jessika E. Trancik and Béla Nagy for helpful comments and discussions.

## Appendix

### A Series expansion approximation to $f$

The standard Maclaurin series expansion for  $(1+x)^\alpha$  is

$$(1+x)^\alpha = \sum_{n=0}^{\infty} \binom{\alpha}{n} x^n \quad \text{for } |x| < 1 \quad (23)$$

$$= 1 + \alpha x + \frac{\alpha(\alpha-1)}{2!} x^2 + \frac{\alpha(\alpha-1)(\alpha-2)}{3!} x^3 + \dots \quad (24)$$

where  $\binom{\alpha}{n}$  are the generalized binomial coefficients

$$\binom{\alpha}{n} = \frac{\alpha(\alpha-1)\dots(\alpha-n+1)}{n!}, \quad \alpha \in \mathbb{C}. \quad (25)$$

Setting  $x$  to  $q^i/z_0^i$  (with  $q^i < z_0^i$ ) and  $\alpha$  to  $-\alpha^i$  gives

$$\left(1 + \frac{q^i}{z_0^i}\right)^{-\alpha^i} = 1 - \alpha^i \left(\frac{q^i}{z_0^i}\right) + \frac{\alpha(\alpha+1)}{2} \left(\frac{q^i}{z_0^i}\right)^2 - \frac{\alpha(\alpha+1)(\alpha+2)}{6} \left(\frac{q^i}{z_0^i}\right)^3 + \dots \quad (26)$$

Then from Eq. (12) we have

$$f = \sum_{i=A,B} c_0^i q^i \left(1 + \frac{q^i}{z_0^i}\right)^{-\alpha^i} e^{(\sigma^i)^2/2} + \lambda (c_0^i q^i)^2 \left(1 + \frac{q^i}{z_0^i}\right)^{-2\alpha^i} e^{(\sigma^i)^2} \left(e^{(\sigma^i)^2} - 1\right),$$

and hence

$$\begin{aligned} f = & \sum_{i=A,B} c_0^i q^i \left(1 - \alpha^i \left(\frac{q^i}{z_0^i}\right) + \frac{\alpha(\alpha+1)}{2} \left(\frac{q^i}{z_0^i}\right)^2 - \dots\right) e^{(\sigma^i)^2/2} \\ & + \lambda (c_0^i q^i)^2 \left(1 - 2\alpha^i \left(\frac{q^i}{z_0^i}\right) + \dots\right) e^{(\sigma^i)^2} \left(e^{(\sigma^i)^2} - 1\right). \end{aligned} \quad (27)$$

Thus if  $q^i$  is much smaller than  $z_0^i$ , then the approximation formed by discarding everything except the zero order terms in the series expansions will be reasonable. This gives the lowest order approximation to  $f$ ,

$$f \approx \sum_{i=A,B} c_0^i q^i e^{(\sigma^i)^2/2} + \lambda (c_0^i q^i)^2 e^{(\sigma^i)^2} (e^{(\sigma^i)^2} - 1), \quad (28)$$

which has the same form as the Markowitz objective function, with  $c_0^i e^{(\sigma^i)^2/2}$  playing the role of expected return and  $(c_0^i)^2 e^{(\sigma^i)^2} (e^{(\sigma^i)^2} - 1)$  the role of the variance. Note that although the series expansion Eq. (26) converges provided  $q^i < z_0^i$ , and is itself reasonably approximated by its zero order term (i.e. 1) provided  $q^i \ll z_0^i$ , the lowest order approximation to  $f$  (Eq. (28)) involves discarding the  $\frac{(q^i)^2}{z_0^i}$  term in the expectation part of Eq. (27), so is reasonable only when  $(q^i)^2 \ll z_0^i$ . Thus the correct condition for the full optimization problem to be Markowitz-like is  $K^2 \ll z_0^i \forall i$ . Higher order approximations to  $f$ , and the corresponding conditions for validity, may be computed similarly. See e.g. Kraus & Litzenberger (1976) for further details on portfolio selection involving preferences over skewness.

## B Analytical points to accompany Section 3.3

### B.1 Onset of diversification

For a given pair of technologies the risk neutral portfolio ( $\lambda = 0$ ) always concentrates production entirely in one technology due to increasing returns, while for sufficiently large risk aversion the portfolio is diversified over both technologies. The value of risk aversion at which the transition between these two regimes occurs (i.e. the onset of diversification) may be found analytically by calculating the intersection of corner and interior solutions. This is done by substituting the relevant boundary condition ( $q^A = 0$  or  $q^A = K$ ) in the first-order condition equation ( $f'(q^A) = 0$ ), and solving for  $\lambda$ .

For example, consider what happens when  $\alpha^B \approx 0.8$  in Fig. 1. For low risk aversion the optimal portfolio is 100% technology  $B$ , but as  $\lambda$  increases there is a value,  $\lambda_{\text{diversification}}$ , at which technology  $A$  first enters the portfolio. Setting  $f'(0) = 0$  (with  $f$  given by Eq. (8)) and solving for  $\lambda$  gives

$$\lambda_{\text{diversification}}(\alpha^B) = \frac{c_0^A e^{(\sigma^A)^2/2} - c_0^B \left( \frac{z_0^B}{z_0^B + K} \right)^{\alpha^B} e^{(\sigma^B)^2/2} \left( 1 - \frac{\alpha^B K}{z_0^B + K} \right)}{2K (c_0^B)^2 \left( \frac{z_0^B}{z_0^B + K} \right)^{2\alpha^B} e^{(\sigma^B)^2} (e^{(\sigma^B)^2} - 1) \left( 1 - \frac{\alpha^B K}{z_0^B + K} \right)} \quad (29)$$

$$= \frac{\mathbb{E}[c_1^A(0)] - \mathbb{E}[c_1^B(K)] \left( 1 - \frac{\alpha^B K}{z_0^B + K} \right)}{2K \text{Var}(c_1^B(K)) \left( 1 - \frac{\alpha^B K}{z_0^B + K} \right)}. \quad (30)$$

Inserting the parameter values from Table 1 in this expression and setting  $\alpha^B = 0.8$  yields the value  $\lambda_{\text{diversification}} = 0.255$ , which is of course consistent with Fig. 1. Note that this expression is independent of  $\alpha^A$ , and only has the interpretation given (onset of diversification) in the specific region of parameter space stated — it does not in itself determine whether  $q^A = 0$  is a global minimum in the first place; that must be verified separately. There is an equivalent expression for the onset of diversification starting from

the 100% technology  $A$  portfolio, computed in precisely the same way. Having found analytical expressions for these two major features of the surface, this just leaves the surface discontinuity itself, which we consider next.

## B.2 Location of the discontinuity

Observe that in Fig. 3 the value of the global minimum varies continuously with  $\alpha^B$ . This follows from the smoothness of the objective function. In fact the global minimum varies continuously with all the underlying parameters, in particular over the entire  $\alpha^B$ - $\lambda$  grid of Fig. 1. Hence the value that  $f$  attains *at* the surface discontinuity in Fig. 1 is the same when approached from both sides, so we have  $f(q_*^A) = f(K - q_*^A)$ . This means that near the discontinuity, portfolios on either side of it are symmetric about a 50% production share. For example, in Figs. 3 and 4 the neighbouring portfolios either side of the discontinuity (i.e. the point at which the optimum switches) are seen to be symmetric about 50%, at approximately 20% and 80% shares. Additionally, in the case where  $q_*^A$  is known exactly (i.e. in the small- $\lambda$ , full specialization region, where  $q_*^A = 0$  or  $K$ ), this insight allows the location of the discontinuity to be calculated analytically.

For example, consider what happens for fixed  $\alpha^B = 0.65$ . As risk aversion increases from zero the optimal portfolio switches from 0 to 100% technology  $A$ . The critical value at which this switch occurs,  $\lambda_{switch}$ , is found by solving  $f(0) = f(K)$  for  $\lambda$ , which gives

$$\begin{aligned} \lambda_{switch} &= \frac{c_0^A \left( \frac{z_0^A}{z_0^A + K} \right)^{\alpha^A} e^{(\sigma^A)^2/2} - c_0^B \left( \frac{z_0^B}{z_0^B + K} \right)^{\alpha^B} e^{(\sigma^B)^2/2}}{K \left[ (c_0^B)^2 \left( \frac{z_0^B}{z_0^B + K} \right)^{2\alpha^B} e^{(\sigma^B)^2} (e^{(\sigma^B)^2} - 1) - (c_0^A)^2 \left( \frac{z_0^A}{z_0^A + K} \right)^{2\alpha^A} e^{(\sigma^A)^2} (e^{(\sigma^A)^2} - 1) \right]} \\ &= \frac{\mathbb{E}[c_1^A(K)] - \mathbb{E}[c_1^B(K)]}{K [\text{Var}(c_1^B(K)) - \text{Var}(c_1^A(K))]} \end{aligned} \quad (31)$$

Again though, note that this expression only has the interpretation given (location of the discontinuity) in the specific region of parameter space stated — it does not in itself determine whether  $q^A = 0, K$  are global minima; this must be verified separately.

The same technique may also be used along the  $\alpha^B$  axis, to compute the value  $\alpha_{switch}^B$  at which the discontinuity occurs, for fixed  $\lambda$ . There is no elementary closed-form solution for this in general, though there is for the special case  $\lambda = 0$ : solving  $f(0) = f(K)$  for  $\alpha^B$ , with  $\lambda = 0$ , gives

$$\alpha_{switch}^B|_{\lambda=0} = \frac{\log \left( \frac{c_0^A}{c_0^B} \right) + \frac{1}{2} [(\sigma^A)^2 - (\sigma^B)^2] + \alpha^A \log \left( \frac{z_0^A}{z_0^A + K} \right)}{\log \left( \frac{z_0^B}{z_0^B + K} \right)}. \quad (32)$$

Inserting the relevant parameters we find  $\alpha_{switch}^B|_{\lambda=0} = 0.596$  in Fig. 1. For other values of  $\lambda$  the solution may be computed numerically. For example, fixing  $\lambda = 0.1$  and solving  $f(0) = f(K)$  for  $\alpha^B$  yields the value  $\alpha_{switch}^B = 0.681$ , which again matches Fig. 1.

## C Demand-risk surface for similar technologies

In order to provide some comparison with the case of the two asymmetrical technologies (Fig. 12), Fig. 15 shows the analogous surface for the two almost identical technologies of




Figure 15: Surface of optimal share of technology  $A$  production as total demand and risk aversion vary, for the two almost identical technologies of Section 3, with  $\alpha^B = 0.65$ .

Section 3, with  $\alpha^B = 0.65$  (so technology  $B$  has slightly higher experience exponent but also slightly higher noise variance). Since the technologies have such similar characteristics, there are very few areas of the  $K$ - $\lambda$  space shown in which complete specialisation is optimal. (Although the  $K$  axis scale is different, the other features roughly match Fig. 12, so the comparison is justified.) In contrast to Fig. 12, the optimal mix is now diversified (except for very small  $K$ ) for any value of  $\lambda$  above around 0.2, and either technology can dominate, depending on  $K$ . Fig. 6 is precisely the  $\lambda = 0.25$  section through this plot, which explains the nearly-Markowitz origin of the “ledge” observed here for very small  $K$ .

## D Two-period objective function

Here we provide expressions for the expectation and variance components of the two-period objective function. We have

$$f(q_1^A, q_2^A) = \mathbb{E} [V(q_1^A, q_2^A)] + \lambda \text{Var} (V(q_1^A, q_2^A)), \quad (33)$$

where

$$V(q_1^A, q_2^A) = \sum_{t=1,2} \sum_{i=A,B} e^{-r(t-1)} c_t^i q_t^i \quad (34)$$

$$= \sum_{i=A,B} c_1^i q_1^i + e^{-r} c_2^i q_2^i \quad (35)$$

is the total system present discounted cost, and first and second period technology costs are given by

$$c_1^i = c_0^i \left( \frac{z_0}{z_1} \right)^{\alpha^i} e^{\eta_1^i} \quad \text{and} \quad c_2^i = c_0^i \left( \frac{z_0}{z_2} \right)^{\alpha^i} e^{\eta_1^i + \eta_2^i} \quad (36)$$

respectively. The noise shocks,  $\eta_1^A, \eta_2^A \sim \mathcal{N}(0, (\sigma^A)^2)$  and  $\eta_1^B, \eta_2^B \sim \mathcal{N}(0, (\sigma^B)^2)$ , are all independent. Now, in the one period problem independence of technologies means there is no covariance term in  $\text{Var}(V)$ , but here, even though technologies are still independent, the second period cost of each technology depends on its own first period cost, so there are nonzero covariance terms to consider. To simplify notation write

$$\bar{c}_t^i = c_0^i \left( \frac{z_0}{z_t} \right)^{\alpha^i} \quad (37)$$

so that

$$c_1^i = \bar{c}_1^i e^{\eta_1^i} \quad \text{and} \quad c_2^i = \bar{c}_2^i e^{\eta_1^i + \eta_2^i}. \quad (38)$$

The following results are then required:

$$\mathbb{E} \left[ e^{\eta_1^i} \right] = e^{(\sigma^i)^2/2} \quad (39)$$

$$\mathbb{E} \left[ e^{\eta_1^i + \eta_2^i} \right] = e^{(\sigma^i)^2} \quad (40)$$

$$\text{Var} \left( e^{\eta_1^i} \right) = e^{(\sigma^i)^2} \left( e^{(\sigma^i)^2} - 1 \right) \quad (41)$$

$$\text{Var} \left( e^{\eta_1^i + \eta_2^i} \right) = e^{2(\sigma^i)^2} \left( e^{2(\sigma^i)^2} - 1 \right) \quad (42)$$

$$\text{Cov} \left( e^{\eta_1^i}, e^{\eta_1^i + \eta_2^i} \right) = e^{3(\sigma^i)^2/2} \left( e^{(\sigma^i)^2} - 1 \right). \quad (43)$$

The expectation component of  $f$  is thus given by

$$\mathbb{E}[V] = \mathbb{E} \left[ \sum_{i=A,B} c_1^i q_1^i + e^{-r} c_2^i q_2^i \right] \quad (44)$$

$$= \sum_{i=A,B} \mathbb{E} \left[ c_1^i q_1^i \right] + \mathbb{E} \left[ e^{-r} c_2^i q_2^i \right] \quad (45)$$

$$= \sum_{i=A,B} \mathbb{E} \left[ \bar{c}_1^i q_1^i e^{\eta_1^i} \right] + \mathbb{E} \left[ e^{-r} \bar{c}_2^i q_2^i e^{\eta_1^i + \eta_2^i} \right] \quad (46)$$

$$= \sum_{i=A,B} \bar{c}_1^i q_1^i e^{(\sigma^i)^2/2} + e^{-r} \bar{c}_2^i q_2^i e^{(\sigma^i)^2}, \quad (47)$$



and the variance component by

$$\text{Var}(V) = \text{Var}\left(\sum_{i=A,B} c_1^i q_1^i + e^{-r} c_2^i q_2^i\right) \quad (48)$$

$$= \sum_{i=A,B} \text{Var}(c_1^i q_1^i + e^{-r} c_2^i q_2^i) \quad (49)$$

$$= \sum_{i=A,B} \text{Var}\left(\bar{c}_1^i q_1^i e^{\eta_1^i} + e^{-r} \bar{c}_2^i q_2^i e^{\eta_1^i + \eta_2^i}\right) \quad (50)$$

$$= \sum_{i=A,B} (\bar{c}_1^i q_1^i)^2 \text{Var}(e^{\eta_1^i}) + (e^{-r} \bar{c}_2^i q_2^i)^2 \text{Var}(e^{\eta_1^i + \eta_2^i}) \\ + 2e^{-r} \bar{c}_1^i \bar{c}_2^i q_1^i q_2^i \text{Cov}(e^{\eta_1^i}, e^{\eta_1^i + \eta_2^i}) \quad (51)$$

$$= \sum_{i=A,B} (\bar{c}_1^i q_1^i)^2 e^{(\sigma^i)^2} (e^{(\sigma^i)^2} - 1) + (e^{-r} \bar{c}_2^i q_2^i)^2 e^{2(\sigma^i)^2} (e^{2(\sigma^i)^2} - 1) \\ + 2e^{-r} \bar{c}_1^i \bar{c}_2^i q_1^i q_2^i e^{3(\sigma^i)^2/2} (e^{(\sigma^i)^2} - 1). \quad (52)$$

## References

- Alberth, S. & Hope, C. (2007), ‘Climate modelling with endogenous technical change: Stochastic learning and optimal greenhouse gas abatement in the PAGE2002 model’, *Energy Policy* **35**(3), 1795–1807.
- Alchian, A. (1963), ‘Reliability of progress curves in airframe production’, *Econometrica* **31**(4), 679–693.
- Almgren, R. & Chriss, N. (2001), ‘Optimal execution of portfolio transactions’, *Journal of Risk* **3**, 5–40.
- Arrow, K. J. (1962), ‘The economic implications of learning by doing’, *The Review of Economic Studies* pp. 155–173.
- Arthur, W. B. (1989), ‘Competing technologies, increasing returns, and lock-in by historical events’, *The Economic Journal* pp. 116–131.
- Atkinson, A. B. & Stiglitz, J. E. (1969), ‘A new view of technological change’, *The Economic Journal* **79**(315), 573–578.
- Barreto, L. & Kypreos, S. (2004), ‘Endogenizing R&D and market experience in the “bottom-up” energy-systems ERIIS model’, *Technovation* **24**(8), 615–629.
- Boyle, P., Garlappi, L., Uppal, R. & Wang, T. (2012), ‘Keynes meets Markowitz: The trade-off between familiarity and diversification’, *Management Science* **58**(2), 253–272.
- Brueckner, J. K. & Raymon, N. (1983), ‘Optimal production with learning by doing’, *Journal of Economic Dynamics and Control* **6**, 127–135.
- Cabral, L. M. & Riordan, M. H. (1994), ‘The learning curve, market dominance, and predatory pricing’, *Econometrica: Journal of the Econometric Society* pp. 1115–1140.

- Cowan, R. (1991), ‘Tortoises and hares: choice among technologies of unknown merit’, *The Economic Journal* pp. 801–814.
- Criqui, P., Mima, S., Menanteau, P. & Kitous, A. (2015), ‘Mitigation strategies and energy technology learning: an assessment with the poles model’, *Technological Forecasting and Social Change* **90**, 119–136.
- Dasgupta, P. & Stiglitz, J. (1988), ‘Learning-by-doing, market structure and industrial and trade policies’, *Oxford Economic Papers* pp. 246–268.
- David, P. A. (1985), ‘Clio and the economics of QWERTY’, *The American Economic Review* **75**(2), 332–337.
- Della Seta, M., Gryglewicz, S. & Kort, P. M. (2012), ‘Optimal investment in learning-curve technologies’, *Journal of Economic Dynamics and Control* **36**(10), 1462–1476.
- Dutton, J. M. & Thomas, A. (1984), ‘Treating progress functions as a managerial opportunity’, *Academy of Management Review* **9**(2), 235–247.
- Edirisinghe, N. C. P. & Patterson, E. I. (2007), ‘Multi-period stochastic portfolio optimization: Block-separable decomposition’, *Annals of Operations Research* **152**(1), 367–394.
- Gârleanu, N. & Pedersen, L. H. (2013), ‘Dynamic trading with predictable returns and transaction costs’, *The Journal of Finance* **68**(6), 2309–2340.
- Gritsevskiy, A. & Nakićenović, N. (2000), ‘Modeling uncertainty of induced technological change’, *Energy policy* **28**(13), 907–921.
- He, H. & Mamaysky, H. (2005), ‘Dynamic trading policies with price impact’, *Journal of Economic Dynamics and Control* **29**(5), 891–930.
- Kraus, A. & Litzenberger, R. H. (1976), ‘Skewness Preference and the Valuation of Risk Assets’, *Journal of Finance* **31**(4), 1085–1100.
- Krey, V. & Riahi, K. (2013), *Risk Hedging Strategies Under Energy System and Climate Policy Uncertainties*, Springer US, Boston, MA, pp. 435–474.
- Kroll, Y., Levy, H. & Markowitz, H. M. (1984), ‘Mean-variance versus direct utility maximization’, *The Journal of Finance* **39**(1), 47–61.
- Lafond, F., Bailey, A. G., Bakker, J. D., Rebois, D., Zadourian, R., McSharry, P. & Farmer, J. D. (2018), ‘How well do experience curves predict technological progress? a method for making distributional forecasts’, *Technological Forecasting and Social Change* **128**, 104–117.
- Majd, S. & Pindyck, R. S. (1989), ‘The learning curve and optimal production under uncertainty’, *The Rand Journal of Economics* **20**(3), 331–343.
- Markowitz, H. (1952), ‘Portfolio selection’, *The Journal of Finance* **7**(1), 77–91.
- Mazzola, J. B. & McCardle, K. F. (1996), ‘A bayesian approach to managing learning-curve uncertainty’, *Management Science* **42**(5), 680–692.

- Mazzola, J. B. & McCardle, K. F. (1997), ‘The stochastic learning curve: Optimal production in the presence of learning-curve uncertainty’, *Operations Research* **45**(3), 440–450.
- McNerney, J., Farmer, J. D. & Trancik, J. E. (2011), ‘Historical costs of coal-fired electricity and implications for the future’, *Energy Policy* **39**(6), 3042–3054.
- Nagy, B., Farmer, J. D., Bui, Q. M. & Trancik, J. E. (2013), ‘Statistical basis for predicting technological progress’, *PLOS ONE* **8**(2), e52669.
- Newbold, P., Pfaffenzeller, S. & Rayner, A. (2005), ‘How well are long-run commodity price series characterized by trend components?’, *Journal of International Development* **17**(4), 479–494.
- Poston, T. & Stewart, I. (2014), *Catastrophe theory and its applications*, Courier Corporation.
- Pulley, L. B. (1981), ‘A general mean-variance approximation to expected utility for short holding periods’, *The Journal of Financial and Quantitative Analysis* **16**(3), 361–373.
- Rosen, S. (1972), ‘Learning by experience as joint production’, *The Quarterly Journal of Economics* pp. 366–382.
- Rosenberg, N. (1982), *Inside the black box: technology and economics*, Cambridge University Press.
- Spence, A. M. (1981), ‘The learning curve and competition’, *The Bell Journal of Economics* pp. 49–70.
- Thompson, P. (2012), ‘The relationship between unit cost and cumulative quantity and the evidence for organizational learning-by-doing’, *The Journal of Economic Perspectives* **26**(3), 203–224.
- Van den Bergh, J. C. (2008), ‘Optimal diversity: increasing returns versus recombinant innovation’, *Journal of Economic Behavior & Organization* **68**(3), 565–580.
- Van Nieuwerburgh, S. & Veldkamp, L. (2010), ‘Information acquisition and under-diversification’, *The Review of Economic Studies* **77**(2), 779–805.
- Wagner, F. (2014), Optimal investment strategy for competing learning technologies: An analytical approach, in ‘Renewables in Future Power Systems’, Springer, pp. 75–129.
- Webster, M., Fisher-Vanden, K., Popp, D. & Santen, N. (2015), Should we give up after Solyndra? Optimal technology R&D portfolios under uncertainty, Technical report, National Bureau of Economic Research.
- Wright, T. P. (1936), ‘Factors affecting the cost of airplanes’, *Journal of the Aeronautical Sciences* **3**(4), 122–128.
- Zeeman, E. (1976), ‘Catastrophe theory’, *Scientific American* **234**, 65–83.
- Zeppini, P. (2015), ‘A discrete choice model of transitions to sustainable technologies’, *Journal of Economic Behavior & Organization* **112**, 187–203.
- Zeppini, P. & Van den Bergh, J. C. (2013), ‘Optimal diversity in investments with recombinant innovation’, *Structural Change and Economic Dynamics* **24**, 141–156.