

Artificial Intelligence, Scientific Discovery, and Commercial Innovation*

Ajay Agrawal

University of Toronto and NBER

John McHale

National University of Ireland, Galway and
Whitaker Institute for Innovation and Societal Change

Alexander Oettl

Georgia Institute of Technology and NBER

July 27, 2019

Preliminary - please do not circulate

Abstract

Motivated by examples of machine learning use in genomics, drug discovery, and materials science, we develop a multi-stage combinatorial model of artificial intelligence (AI) -aided innovation. The innovator can utilize AI as a tool for drawing on existing knowledge in the form of data on past successes and failures to produce a prediction model - or map - of the combinatorial search space. Modeling innovation as a multi-stage search process, we explore how improvements in AI could affect the productivity of the discovery pipeline in combinatorial-type research problems by allowing improved prioritization of the leads that flow through that pipeline. Furthermore, we show how enhanced prediction can increase or decrease the demand for downstream testing, depending on the type of innovation. Finally, we examine the role of data generation as an alternative source of spillovers in sustaining economic growth.

*We gratefully acknowledge financial support from Science Foundation Ireland, Social Sciences Research Council of Canada, Centre for Innovation and Entrepreneurship at Rotman School of Management, and the Whitaker Institute for Innovation and Societal Development, NUI Galway. All errors are our own. © 2019 Ajay Agrawal, John McHale, and Alexander Oettl.

1 Introduction

In December 2018, Google DeepMind’s AlphaFold won the 13th round of the CASP¹ protein folding competition. They used a type of AI, machine learning using deep neural networks, to predict the 3D structure of target proteins based on their amino acid sequences. The “protein folding problem” is challenging due to the vast number of potential shapes that a protein could take for a particular amino acid sequence. Knowing the shape of proteins is, among other uses, critical for identifying targets for drugs to bind to in order to produce therapeutic effects. Such applications of AI are enhancing discovery in a domain that previously yielded a disappointing output of new drugs despite apparent rapid scientific advance.

In April 2017, Atomwise, an AI-based drug discovery company, announced the Artificial Intelligence Molecular Screen (AIMS) Awards. Under this program, university-based awardees submit a target protein and then receive 72 candidate compounds in assays for testing. The candidate compounds are identified for the awardee by Atomwise, using their proprietary AI-based system to run a virtual screen on a vast library of chemical compounds. Atomwise benefits from providing this free service to awardees by gaining access to training data to further improve its AI-based prediction algorithms.

These are just two examples of the use of AI in scientific discovery and innovation. AI tools are increasingly being applied by innovators over large and complex search spaces, promising to improve the productivity of the innovation process in various domains. The new tools also raise new challenges for policy makers, university administrators, and managers of innovation-intensive businesses as they seek to organize the innovation process for societal and institutional benefit.

We develop a model of an AI-aided innovation process. The model is motivated by four main ideas. First, innovation can be viewed as search over a potentially vast combinatorial search space. Second, the uncertainty innovators face as to the location of valuable combinations in this space can be partly reduced by having a prediction model – or map – of the underlying search landscape. Third, the output of the prediction model can usefully be summarized in the form of a ranking function that allows prioritization in the context of a costly multi-stage innovation search process.

¹Community Wide Experiment on the Critical Assessment of Techniques for Protein Structure Prediction (CASP).

And fourth, advances in AI are a potential source of improvement in the performance of these predictive maps.

Scholars have long modeled innovation as a process of combining existing ideas to produce new ideas (Usher, 1929; Schumpeter, 1939; Nelson and Winter, 1982; Weitzman, 1998; Fleming and Sorenson, 2004; Arthur, 2009; Agrawal et al., 2016; Teodoridis, 2018). The notion of creating new knowledge by combining existing pieces of knowledge in new ways has received particular attention where the mapping is highly complex. Example domains include biotechnology, molecular and materials science, and particle physics. Recent advances in AI – and in particular advances in machine learning techniques that employ neural networks, such as deep learning – have led to optimism that these advances will produce a new general purpose technology (GPT) for discovery (Agrawal et al., 2019a; Cockburn et al., 2019). AI-based tools are already enhancing innovation in fields such as bioinformatics, cheminformatics, and genomics, although some skepticism remains about their ultimate impact on R&D productivity.

We build on the classic innovation function approach to economic growth by developing a model of AI-aided innovation. The standard approach models innovation as a function of research effort and existing knowledge stocks (Romer, 1990; Grossman and Helpman, 1991; Aghion and Howitt, 1992; Jones, 1995). In our model, an innovator’s existing knowledge enters the discovery process in two ways. First, their knowledge determines the search space of potential combinations. Second, they can use additional knowledge – data on past successes and failures – to develop a prediction model for the “fitness landscape” that maps combinations to the probability of success of those combinations. In essence, the model of the fitness landscape aids their search in the context of a multi-stage search process by helping to predict which combinations out of possibly billions have the greatest likelihood of success. We thus treat a key part of the discovery process as a prediction problem that can benefit from recent advances in AI.

A summary measure of the output of the prediction model is a ranking function that shows how the probability of success declines as we move from better- to worse-ranked combinations. This function plays a key role in the prioritization process for a multi-stage discovery pipeline where the later stages – typically testing – are costly, and access to improved technologies for prediction can

better prioritize the use of costly R&D resources.

The paper is related to a number of literatures. First, our paper is inspired by a growing literature that explores the use of AI in scientific discovery and innovation. Rapid advances in hardware, algorithms, and data availability have driven this growth. For a sampling of recent reviews see Chen et al. (2018) [drug discovery], Angermueller et al. (2016) [computational biology], Wainberg et al. (2018) [genomics], Goh et al. (2017) [computational chemistry], and Butler et al. (2018) [materials science]. The success of AI in providing predictive models for highly complex combinatorial spaces explains the rising interest in AI as a GPT for discovery. This potential – together with the power of viewing innovation as a prediction-model-aided combinatorial search process – motivates the paper.

Second, we draw on and extend the innovation production function that is at the core of endogenous growth theory (see, e.g., Jones, 1995). Central to the innovation production function approach is the idea that existing knowledge is an input into the production of new knowledge – the “standing on the shoulders of giants” or spillover effect. Typically, papers in this literature do not explicitly adopt a combinatorial view of the process through which existing knowledge is turned into new knowledge (although see Romer, 1993; Weitzman, 1998, for exceptions). However, the idea of search over combinatorial spaces is central to our approach. We also highlight a second source of potential spillovers: data on past successes and failures. These data are used to develop improved prediction models and thus can be a source of growth-sustaining productivity improvement in the innovation search process.

Third, our paper draws on literature in economics that applies the ideas of fitness landscapes to the study of innovation. Sewall Wright (Wright, 1932) first introduced the fitness landscape concept in evolutionary biology and Stuart Kauffman (see Kauffman (1993)) extensively developed it. As used in economics and management science, this work is typically situated in evolutionary economics and builds on the pioneering work of Nelson and Winter (1982). Important contributions include Levinthal (1997), Gavetti and Levinthal (2000), Kauffman et al. (2000), Rivkin (2000), Fleming (2001), and Fleming and Sorenson (2004). Fleming and Sorenson (2004) introduced the idea of science as a map to aid technological search, an idea that is central to our approach. Drawing

on the evolutionary approach, these papers model innovation (or imitation) as a “walk” or “hill climb” towards a local optimum on the fitness landscape. Innovators typically search one-mutant neighbors and adopt fitter variants until a local optimum is reached, although “long-jumps,” in which innovators jump longer distances across the landscape, are also studied as a process of exploratory search. The evolutionary approach has proved extremely fruitful and provides rich dynamics for the search process. However, to better connect with the innovation function approach that has been standard in the endogenous growth literature, we do not use an evolutionary approach in this paper. Instead, in our model, innovators use (imperfect) knowledge of the fitness landscape (the prediction model) to identify a promising subset of potential combinations followed by testing of that subset.

Fourth, we draw on the literature on optimal search where information is imperfect and search is costly. This literature originated with Stigler (1961) with influential developments in McCall (1965). We draw in particular on a special case of the “Pandoras box” model developed in Weitzman (1979). The sequential search problem examined by Weitzman involves boxes that vary in the distribution of potential outcomes, search costs and the time that elapses before the value of the box is revealed. He shows that a “reservation price” can be attached to each box. Ranking the boxes in descending order of reservation prices, the optimal search rule – “Pandoras rule” – is to continue down the ranking until the maximum value obtained is greater than the reservation prices of all remaining unopened boxes.

We utilise a special case of the Weitzman model where there are just two outcomes that can be revealed in the test of a given combination – success or failure – and there is an unbiased estimate of the probability of success available to a risk-neutral innovator from the prediction model. We also assume that costs are the same for all tests and that there is no time discounting. This allows the innovator to prioritize solely based on the ranking of combinations given by the prediction model and we make a strong assumption about the functional form of the ranking function. Using this tractable search set up, we examine two cases of simultaneous (or parallel) search where (1) the innovator seeks to discover all combinations with an expected net value greater than or equal to zero and (2) the innovator seeks a single success but must choose in advance which combinations are to

go for testing.² Drawing on the real options literature, we also extend our two-stage search process to a more general multi-stage search process where there is an option to abandon an advancing combination at the end of each stage (see, e.g., Roberts and Weitzman, 1981; Dixit et al., 1994).

Finally, our paper draws on contributions to the emerging literature on the economics of artificial intelligence (AI). A key breakthrough in AI has followed from the shift from rules-based systems to a statistical approach that emphasizes prediction (see, e.g., Athey (2017); Mullainathan and Spiess (2017); Agrawal et al. (2019b); Taddy (2019)). As emphasized by Arrow (1962), uncertainty is a pervasive feature of the innovation process and hence the value of prediction technologies that help reduce it. We view AI as a GPT for prediction and make extensive use of this idea in our paper.³

We organize the remainder of the paper as follows. Section 2 briefly illustrates how machine learning is altering the discovery process in genomics, drug discovery, and materials science. Section 3 sets out the basic conceptual building blocks of a search-based innovation model and relates it to the innovation production function and the idea of a fitness landscape. Section 4 then develops a simple two-stage example of “search with a map” in which the first stage is the development of a prediction model (the output of which is captured by a logistic ranking function) combined with identification of a probability threshold for determinative testing and the second stage is the actual testing of all combinations with a probability of success at or above the threshold. Section 5 introduces machine learning as a method of generating improved prediction models of the fitness landscape. Section 6 extends the discussion to a multi-stage setting to allow for a sequential refinement of predictions and an option to abandon before engaging in expensive testing of candidate contributions. Section 7 considers a number of issues that have particular salience in the context of a multi-stage discovery process: bottlenecks in the innovation process, autonomous innovation processes that minimize those bottlenecks, and the implications of AI-aided innovation on the demands for specific forms of R&D labor. Section 8 briefly explores the role of data spillovers as a

²Chade and Smith (2006) provide a more general treatment of the simultaneous search problem.

³Another important strand of the economics of AI literature has focused on the effects of AI on the demands for different types of skill. Researchers in this area have used the idea of a task-based production function to allow for the possibility that the introduction of AI (and other new technologies such as robotics) could lower the employment and wages of certain types of workers depending on the tasks they perform (Acemoglu and Autor, 2011; Autor, 2015; Acemoglu and Restrepo, 2019a, 2018). Such labor demand effects could take place for knowledge production tasks as well (Aghion et al., 2019).

second source of spillovers (after ideas) in the innovation production function. Section 9 concludes with a recap of the main ideas and a discussion of the possible policy and managerial implications of AI-aided scientific discovery and innovation.

2 Innovation as Search over Complex Spaces: Motivating Examples

To motivate our modeling approach, we illustrate discovery challenges that involve search over complex combinatorial spaces, challenges for which AI-based prediction models appear to be fruitful. Often, an innovation, such as a new DNA sequence or a chemical molecule, has combination-specific properties or activities, such as a relationship to an intermediating cell variable or a ligand-protein binding affinity. The discovery processes are typically involve multiple stages (e.g., predictive screening, synthesis, testing, etc.). We are interested in how researchers use AI to screen candidate combinations and rank them to optimize further exploration along the discovery pipeline.

Our first example is from genomic medicine. A genome is a complex set of instructions for building an organism. The building process can be viewed as a combinatorial problem: for a given gene with a DNA sequence of length L , the number of possible combinations of chemical bases (C, G, A and T) is (ignoring redundancies) 4^L . Genomic medicine exploits the relationship between DNA sequences and the risk of various diseases. A central idea is that of gene expression, the process by which the information in the gene is first transcribed to make messenger RNA (mRNA) and then the mRNA is translated to make a protein (Leung et al., 2016).⁴ Scientists can utilize predictive models for various stages of this process. An encompassing approach is to model the relationship between DNA sequences and disease outcomes. However, given the complexity of the process, scientists are making significant progress by developing predictive models of the relationship between DNA sequences and various intermediating “cell variables” (Leung et al., 2016). These cell variables can provide potential targets for therapeutic interventions. New measurement methods can give high-throughput data on various cell variables, and advances in AI are allowing scientists

⁴We can view a protein itself as a combination of amino assets. For an amino asset polymer of length L^* , given the 20 common amino assets, there are 20^{L^*} possible sequences of that polymer.

to take advantage of these abundant data to develop predictive models of their functioning.

Our second example comes from small molecule drug discovery. Small molecule drugs are a mainstay of therapeutic medicine. Scholars estimate the space of potential small organic molecules to contain more than 10^{60} possible structures (Virshup et al., 2013). These chemical combinations (or ligands) interact with targets (e.g., a protein) that regulate biochemical processes within the body. Disease can occur when these targets malfunction. By binding to a malfunctioning target protein, the ligand may be able to alter its adverse bioactivity. As combinations of chemical elements, these small molecules can take a vast number of forms leading to a massive combinatorial search space even for a single protein target. Taking the target (or “lock”) as given, the challenge is then to identify a ligand (or “key”) that binds effectively and leads to the desired therapeutic effect. This screening challenge can be aided by AI models of binding efficacy (e.g., Chen et al. (2018)). Some recent approaches take advantage of knowledge of the three-dimensional structures of both the target proteins and the small molecule ligands. Machine learning models such as convolutional neural networks (CNNs) – initially applied to tasks such as image recognition – are utilizing this information to improve predictions of bioactivity for drug discovery applications (see, e.g., Wallach et al. (2015); Gomes et al. (2017)).

Our final example comes from materials discovery. As with drug discovery, the space of potential molecules is vast. Researchers have used computational chemistry to virtually screen for the properties of molecules, including methods such as quantum chemistry and molecular mechanics. However, the computational costs of such simulation methods can be prohibitive, leading to increasing interest in statistical approaches such as AI to prioritize molecules for simulation- or experiment-based characterization (see, e.g., Pyzer-Knapp et al., 2015). For example, input data on molecular descriptors and output data on molecular properties could be used to develop a machine-learning-based predictive model of a large chemical space that would otherwise be prohibitively costly through computational and experimental methods.

Materials discovery – for instance, the search for new materials for clean energy technologies or medical devices – is well characterized by multi-stage search over a combinatorial search space. The tasks involved in the discovery of new materials include predictive screening of potential molecules,

making (or synthesizing) those molecules, testing the molecules using high-throughput methods, and characterizing their properties. Among the innovation challenges to which AI is being applied are the development of new catalysts to convert earth-abundant molecules (such as CO₂, H₂O and N₂) into fuels and chemicals, new photovoltaic and thermoelectric materials, and new forms of batteries for energy storage (Tabor et al., 2018). However, researchers have concerns that bottlenecks in the discovery process severely slow the flow of new discoveries.

Motivated by advances in both AI and robotics, scientists are increasingly interested in “autonomous” (or “self-driving”) discovery systems, which form a closed loop between predictive screening and characterization. Looking in particular at the case of new materials for clean energy, Aspuru-Guzik and Persson (2018, p. 23) set out the challenge of autonomous materials design as combining: (i) AI-based predictions of new materials and their properties, (ii) autonomous robotic systems for synthesis and experimental data collection, (iii) data analysis such as feature extraction, (iv) AI-based classification and regression of results, and (v) decision modules to drive optimal experimental design for subsequent experimental iterations. For autonomous discovery and development, these elements need to be integrated into a closed-loop platform for designing and performing experiments.

As a GPT for discovery, the impact of AI on such an autonomous system may be limited due to bottlenecks in the system. However, as AI and complementary technologies, such as robotics, improve and become better integrated, scientists anticipate that the productivity of the discovery pipeline will be significantly enhanced (Tabor et al., 2018).

3 A Combinatorial Model of Innovation: Building Blocks

Our examples reflect a wide range of types of input data, output data, and algorithms to produce prediction models. In this section we put these details aside and set out the conceptual building blocks of a highly stylized model of the innovation process as search over a vast combinatorial search space.

Our starting point is the innovation production function that is central in much of the endogenous growth literature. Innovation (I) is modeled as new ideas (\dot{A}) and new ideas are a function

of the stock of existing ideas (A) : $I = \dot{A} = F(A)$. This captures the basic spillover effect, whereby new ideas spillover, providing the feedstock for the production of yet further ideas (Romer, 1990; Jones, 1995).

One natural interpretation is that innovation is a combinatorial process such that each idea is a *combination* of other ideas (see Romer, 1993; Weitzman, 1998; Arthur, 2009). We can represent an idea as a string. In the simplest case, the string reflects whether each idea is present or not in the combination, with a 1 when that idea is present and 0 when it is not. More generally, each of the ideas in a string can have multiple states, with the set of M states denoted as its alphabet. This string could represent a DNA sequence, amino acid sequence, a sequence of molecular descriptors or fingerprints, etc. We can then represent the string underlying a given idea (i.e., combination) by the particular states of the ideas that comprise that combination.

For a given innovator, their *ideas access* is their initial knowledge stock. This knowledge access determines their *combinatorial search space* – that is, the set of all possible combinations of the ideas in their knowledge stock that can be formed. In the binary case ($M = 2$), for an innovator with a stock of A ideas, the number of all possible combinations that can be formed from these ideas is 2^A .⁵

Innovation output can now be thought of as a function of the potential combinations: $I = \dot{A} = G(2^A)$. More precisely, we think of innovation as resulting from search over the set of potential combinations.

In conceptualizing this search space we make use of the idea of a *fitness landscape* (Kauffman, 1993). For a given fitness landscape, we associate each potential combination with a scalar that reflects its (fitness) value according to some particular property of interest. Value may be multidimensional, so there can be multiple landscapes, each one associated with a particular property of interest. However, we will typically assume a single dimension of technological fitness for convenience.

Central to the fitness landscape is a measure of distance. The distance between any two strings (or combinations) is the number of states that differ between those strings (or Hamming distance).

⁵More generally, if each idea can take any one of M states, the total number of possible combinations is M^N .

For a given string, the 1-neighbor strings are all those strings that differ by just one state. The d -neighbor strings are all those strings that differ in exactly d of the states.

The correlation structure of the landscape determines how correlated the values of the combinations are across different distances. Low correlations between the values of combinations outside close neighborhoods are associated with more “rugged” fitness landscapes.

The mapping from a combination to its scalar value can be viewed as an index function. For a combination to be successful – i.e., lead to an innovation – we assume its index value must be equal to or greater than some threshold. We can thus recast our landscape in a simplified form so that combinations with a value at or above that threshold have a value of 1 and combinations with a value below the threshold have a value of 0.

We now introduce a second form of knowledge that is available to the innovator: *data*. In addition to ideas – the stock of which determines the combinatorial search space – the innovator has knowledge of previous successes (i.e., ideas with a value of 1) and previous failures (i.e., ideas with a value of 0).

These data can be used by the innovator to develop a fitness landscape in the form of prediction model that outputs the probability of success of any potential combination. One useful summary of the output of the prediction model is a ranking function that ranks the probability of success of each combination in ascending order.

We next develop a simple two-task example of the innovation process that assumes a specific functional form for this ranking function and examine the implications of access to an improved prediction technology for the innovation decision-making process and the expected value of innovation.

4 Searching with a Map: A Two-Stage Example

We consider a simple two-stage example. In the first stage, the innovator develops a model for predicting the value of each possible combination. In the second stage, the innovator tests those combinations that have a predicted value that exceed a threshold. Testing is costly. In the context of this two-stage example, we explore how access to an improved prediction model affects the

innovation process and its outcomes.

In the first stage, our innovator is faced with the task of searching over a possibly vast landscape of $2^A - D$ potential combinations, where A is the number of ideas that the innovator has available to combine into new ideas and D is the number of observations on prior successes and failures. A higher value of D reduces the number of combinations available to be discovered, but also provides “training data” for developing a prediction model for successful new combinations.

This search is characterized by uncertainty over the location of the valuable undiscovered combinations on the landscape, which we assume to be small in number relative to the size of the space. However, our innovator also has access to the prediction model to help detect the location of the undiscovered successes. This prediction model could be provided by theory, simulation, parametric statistical data modeling, AI, or even educated guesses – although all methods will likely rely to some degree on observations of prior successes and failures.

Again, innovation consists of two stages. Stage 1 (prediction) is the development (or “training”) of the prediction model.⁶ We conveniently represent the output of the prediction model as a ranking function that maps the rank (1 for the top ranked combination to $2^A - D$ for the bottom ranked combination) to the model-based probability that the combination is a success.

Stage 2 (testing) involves the conducting of the test, where we assume that testing is a regulatory requirement that cannot be bypassed by taking a promising combination straight to market. The cost of a test is c .

Combinations are ultimately either successes (1) or failures (0), but without conducting the determinative test the innovator is uncertain as to which and must form an expectation of the probability of success prior to testing. If a successful innovation is brought to market after testing it results in a payoff to the innovator of π .

We seek a decision rule at Stage 1 to determine which combinations to send for testing. In

⁶We assume that the production of the prediction model precedes the identification of the probability threshold. As we will see, the probability threshold is actually chosen independently of the prediction model. It is thus possible that the threshold is chosen first, i.e., before the development of the prediction model. This could lead to a better performing model in terms of discrimination between successes and failures if it is important, say, that the model is performing well around the threshold. This would also allow for standard classification-based performance measures (e.g., the accuracy rate) to be used in selecting the prediction model. The qualitative results from the model are not affected by this alternative timing.

determining this decision rule, we assume that the innovator is risk neutral and has the objective to maximise the total expected value (net of the cost of testing) of innovation.

There is a “ground truth” only imperfectly known to the innovator as to the location of undiscovered successes on the landscape. There are G successes in total so that successes as a share of potential combinations is $G/(2^A - D)$, where G is a natural measure of the fecundity of the search space. A graphical representation of this ground truth is the unit step function shown in Figure 1. Denoting the known probability of success of the r^{th} ranked potential combination as p_r , the unit step function is:

$$\begin{aligned} p_r &= 1 && \text{for } r \leq G \\ &= 0 && \text{for } r > G. \end{aligned} \tag{1}$$

Figure 1 shows both the case where there is perfect ability to discriminate between successes and failures (where the internal rankings within the subsets of both successes and failure is arbitrary) and also the case where there is no ability to discriminate between successes and failures. In the latter case, the probability of success is simply the probability of finding a success as a result of a random draw from the search space, i.e., $G/(2^A - D)$.

We seek a functional form for the ranking function such that the probability of discovering a success is equal to $G/(2^A - D)$ when the prediction model has zero discriminating power and approaches the ground truth as the model approaches perfect discrimination. The following logistic decay function has the property that the ground truth will be approached as $b \rightarrow \infty$:

$$p_r = \frac{1}{1 + Ke^{b(r-G)}}, \tag{2}$$

where $b \geq 0$. The value of K can be chosen so that the probability of success is equal to $G/(2^A - D)$ when the prediction model has zero discriminating power (i.e., $b = 0$):

$$\begin{aligned}
p_r &= \frac{1}{1+K} = \frac{G}{2^A - D} \\
\Rightarrow K &= \frac{2^A - D - G}{G},
\end{aligned}
\tag{3}$$

We chose this (appropriately restricted) logistic function as an analytically convenient representation of the ranking function. Furthermore, we assume that b monotonically increases with the performance of the prediction model. As shown in Figure 2, the shape of ranking function curve varies from a horizontal line at $G/(2^A - D)$ when $b = 0$, and converges to the ground truth unit step function as $b \rightarrow \infty$. Thus, the performance of the prediction model is controlled by a single parameter. Increases in b cause the ranking function curve to rotate in a clockwise direction around the point $(G, G/(2^A - D))$.⁷

We now consider two cases of two-task innovation. In Case 1, the value of an innovation is independent of which other valuable innovations are discovered, and our risk-neutral innovator seeks to discover all valuable combinations with a positive expected value net of the cost of testing. In Case 2, the innovator has a single target – say a small molecule drug or material with a particular property. The value to our innovator of finding additional successful combinations once one success is achieved is zero. We consider each case in turn, focusing on how the prediction model affects the decision to select combinations for testing and ultimately the expected total net value of innovation.

⁷Numerous measures of performance exist for binary dependent variable models. Ideally, the measures of performance would be applied to a held back test (or validation sample) given the risk of overfitting, especially for models that allow for highly flexible functional forms. One particularly intuitive measure for evaluating predictive performance both in and out of the estimation sample is the Tjur Coefficient of Discrimination (Tjur, 2009): $E[p(\text{success}|\text{success})] - E[p(\text{failure}|\text{success})]$. This coefficient varies from 0 to 1, with a coefficient of 1 indicating a perfectly discriminating model with the mean probability of success given the combination is an actual success equal to 1 and the mean probability of success given that an actual success equal to 0. The area under the receiver operating characteristic curve (AUROC) is also a widely used measure of the performance of a binary dependent variable model that can be usefully related to the accuracy of the probability of success rankings: the AUROC can be interpreted as the probability that a randomly chosen actual success is ranked better (a lower number given our model) than a randomly chosen actual failure. In general, there is of course no reason that the probability rankings produced by a prediction model will follow a logistic curve. However, it is intuitive that a better performing model will tend to increase the estimated probabilities associated with better ranked combinations and decrease the estimated probabilities associated with poorly ranked combinations. In our model, an increase in the parameter b brings about the required clockwise tilt in the ranking function curve. This fact combined with the analytical convenience of the logistic functional form leads us to treat b as a useful parameter to control the performance of the prediction model.

Case 1: Innovator searches for all profitable innovations

To evaluate the impact of an improved prediction technology on expected value, we cast the testing decision in terms of a comparison of the expected marginal gross value and marginal cost of testing. We denote the optimal number of combinations to send for testing as r^* . Assuming an ordering of tests in terms of decreasing expected marginal gross value of the innovation, the expected marginal gross value of the r^{th} test is:

$$MV_r^e = p_r \pi. \quad (4)$$

the marginal cost of the r^{th} test is simply:

$$MC_r = c. \quad (5)$$

The expected marginal gross value and marginal cost curves are illustrated in Figure 3a. The subset of combinations that will be sent for testing comprises those combinations with an expected marginal gross value greater than marginal cost. If we make the additional simplifying assumption that at the last combination sent for testing the marginal value and marginal cost are equal, then the probability threshold is given by:

$$p_{r^*} = \frac{c}{\pi}. \quad (6)$$

Using equation (2) – our logistic ranking function – allows us to solve for the optional number of tests. All tests where the expected marginal gross value is greater than or equal to the marginal cost will be conducted:

$$\begin{aligned} \frac{1}{1 + \left(\frac{2^A - D - G}{G}\right) e^{b(r^* - G)}} &= \frac{c}{\pi} \\ \Rightarrow r^* &= G - \frac{\ln\left(\frac{2^A - D - G}{G}\right) - \ln\left(\frac{\pi}{c} - 1\right)}{b}. \end{aligned} \quad (7)$$

More generally, the probability threshold (and associated r^*) will be the lowest probability in the ranking at which $MV^e \geq MC$. Expected total net value is then:

$$\begin{aligned}
 V^e &= -cr^* + \pi \sum_{r=1}^{r^*} p_r \\
 &= -c \left(G - \frac{1}{b} \left[\ln \left(\frac{2^A - D - G}{G} \right) - \ln \left(\frac{\pi}{c} - 1 \right) \right] \right) + \pi \sum_{r=1}^{r^*} \left(\frac{1}{1 + \left(\frac{2^A - D - G}{G} \right) e^{b(r-g)}} \right). \quad (8)
 \end{aligned}$$

In Figure 3b we illustrate the impact of an improvement in the prediction model (i.e. an increase in b) on the number of combinations that will be sent for testing. Provided $c > \pi(G/2^A - D)$, the MC curve will intersect the MV^e curves above the crossing point of the MV^e curves at $G/(2^A - D)$. We assume that this condition holds, recalling that the number of undiscovered combinations is assumed to be small compared to the size of the combinatorial search space. The number of combinations sent for testing will be a non-decreasing function of b and strictly increasing for a large enough increase in b for it to be optimal to send at least one additional combination for testing.

Of most interest is what happens to the expected total net value of innovation, V^e . At the original optimal number of tests, r^* , the marginal expected gross value will be greater with the new (higher b) prediction model than with the original model (see Figure 3b). It follows that V^e must be higher even if the number of combinations going for testing did not change. However, any increase in the optimal number of tests from r^* to r^{**} as a result of the improved prediction model will lead to a further increase in V^e . The overall increase in V^e is shown in Figure 3b.

Case 2: Innovator searches for a single target innovation

Case 1 made the strong assumption that the value of alternative innovations are independent and our risk neutral innovator seeks all innovations that yield a positive expected net value. However, in many innovation search problems the innovator may be looking for a *single* combination that meets a particular target such as a small molecule drug that binds with a target protein to improve its functioning or a material for a battery with a desired property. Bringing multiple innovations

to market that achieve the same target will be wasteful to the extent that these innovations are substitutes. We thus examine the opposite extreme of search for a single successful target innovation.

In the classic (“Pandoras box”) search model of Weitzman (1979), the decision maker searches for the single best alternative by deciding on the order to open boxes (i.e., test) and when to stop the search. In Weitzman’s model, the decision maker knows ex ante the probability distribution of outcomes, faces boxes with different opening costs and different time lengths before the contents of any opened box is revealed. Weitzman shows that each box can be assigned a reservation price that depends on the unique characteristics of that box. He derives “Pandoras rule” as the optimal search strategy: continue to open boxes in descending order of the reservation price until the value of the best outcome achieved is greater than the reservation prices of all remaining unopened boxes.

However, Weitzman considers a special case that matches our basic setup: the outcome for a given box is either a success or a failure with a given probability of success for each box, the cost of opening a box is the same for all boxes and the length of time before the outcome is revealed is irrelevant as there is no time discounting. In that case, it is optimal to open all boxes with an expected net value greater than or equal to zero in descending order of the probability of success, and to stop the search when a success is achieved. The decision problem examined here has one important difference: the decision maker must choose all the combinations to be tested (or boxes to be opened) before any of the boxes can be opened. It is thus an example of simultaneous (or parallel) search rather than sequential search. Therefore in deciding on whether to search an additional box the decision maker must take into account that a success may have already been achieved from one of the combinations previously selected for testing.

The difference from Case 1 comes in the position and shape of the expected marginal gross value curve: as our innovator evaluates the expected marginal gross value of an additional test, they must consider the probability that a combination meeting the target would have been discovered with the existing tests. The expected marginal gross value curve is now:

$$\begin{aligned}
MV_r &= \pi p_r && \text{for } r = 1 \\
&= \pi \left(\prod_{j=1}^{r-1} (1 - p_{r-j}) \right) p_r && \text{for } r = 2, 3, \dots, 2^A - D.
\end{aligned} \tag{9}$$

For the marginal test, r , the term inside the first round brackets gives the probability that the desired target combination will not have been discovered in the previous $r - 1$ tests. For example, for the 3^{rd} test the probability that the target will not have been discovered in tests 1 and 2 is $(1 - p_1)(1 - p_2)$. The probability $(1 - p_1)(1 - p_2)p_3$ is thus the probability that target will be found in the third test and not before.⁸

The marginal cost of a test is again given by equation (5). As illustrated in Figure 4a, the optimal number of combinations to send for testing is identified by moving down the probability ranking and testing all combinations for which the expected marginal gross value is greater than or equal to the marginal cost. This gives the optimal number of tests, r^* , and the expected total net value is:

$$V^e = -cr^* + \pi \sum_{t=r}^{r^*} \left(\left(\prod_{j=1}^{r-1} (1 - p_{r-j}) \right) p_r \right). \tag{10}$$

Figures 4b and 4c show the impact on the optimal number of tests when the innovator gets access to an improved prediction model, again captured by an increase in b . As can be seen from the figures, the marginal value curves (pre and post improvement) will cross at a certain point in the ranking. The impact of an improvement in the prediction model will depend on whether marginal cost, c , is above or below marginal value at this crossover point. Figure 4b shows a case where the crossover point occurs below c . In this case the optimal number of tests will increase. Figure 4c shows a case where the crossover occurs above c and the optimal number of tests will

⁸One apparent inconsistency in this set up is that only a single innovation that meets the target has value in the market and the implicit assumption that there is data available on prior successes, which (along with prior failures) allow the training of the prediction model. The presence of prior successes would indicate that an additional success would have no value. One solution is to assume that a new success has to have multiple specific properties and a combination with all those properties has yet to be discovered. The data on successes and failures could be related to individual properties. The relevant ranking curve would then be for the probability that a given combination has all the required properties.

decrease.

What is the effect of an increase in b on V^e ? Provided p_r is greater than $G/(2^A - D)$, we can show that V^e increases regardless of whether the optimal number of tests rises or falls. To see why, note that for a higher value of b the probability of success is higher for each position in the ranking. Comparing the probability that a success will have been found at the r^{th} ranked combination with and without the improvement in the prediction (where a prime indicates the post-improvement probability at a given rank):

$$1 - \prod_{j=1}^r (1 - p'_{r-j}) > 1 - \prod_{j=1}^r (1 - p_{r-j}). \quad (11)$$

Given the improved prediction model, the probabilities on the left-hand-side all greater than the probabilities on the right-hand-side at a given rank. Thus the cumulative probability that a success will have been discovered is also higher for every position in the ranking. In Figure 4b, V^e is thus higher with the higher b at the original optimal number of tests, r^* . The move to the higher optimal number of tests r^{**} increases V^e still further. A similar situation arises in Figure 4c even though the optimal number of tests now falls rather than rises. At the original number of tests V^e is again higher with the higher b . The move to the optimal number of tests (fewer tests in this situation) further increases this gain relative to the situation with the inferior prediction model (i.e. lower b).

Summarizing the analyses of Case 1 and Case 2, although the impact of an improved prediction model on the number of combinations optimally advancing to testing depends on the precise nature of the two-stage search captured in each case, a common result is that an improved prediction model can be shown to increase the expected total net value of innovation. We return to consider a more general multi-stage process in Section 6.

5 Enhanced Prediction: Advances in AI as a Shock to Innovation

We next introduce AI as the source of the shock to the two-stage innovation process. Our generic workflow in the two-task model is prediction followed by testing. Prediction is based on a model

of the fitness landscape. As captured in Figure 5, there are a number of ways to generate such prediction models for the landscape. A useful distinction is between theory- and simulation-based approaches on one hand and data-based approaches on the other. For reasons we explain below, we treat AI as a *subset* of the data-based approach. While data is obviously also central in the testing phase (say through the use of controlled or natural experiments), the focus here is on how data is used to generate prediction models prior to the testing task.

In order to better delineate the role of AI as a Stage 1 prediction tool, we first briefly underline the role of theory in generating predictions. Of course, theory is the means of generating testable predictions in Popper’s classic account of the scientific method Popper (1959). However, we take it that the central requirement for science is that predictions (or conjectures) can be tested, and leave open the source of those predictions to include data-based approaches. Theory, of course, remains a major source of predictions. In theoretical chemistry, for example, the Schrödinger equation remains central to the prediction of molecular properties. Given the complexity of predictions beyond extremely small molecules, various approximations such as the Born-Oppenheimer approximation or Hartree-Fock theory are used. However, even with these approximations the complexity of the calculations requires the use of (typically costly) computer simulations.

The second method of prediction prior to testing is data-based prediction (or hypothesis) generation. We think of AI as a subset of data-based prediction – a subset for which there has been recent rapid progress. In distinguishing this subset, we find it useful to make use of Leo Breiman’s 2001 contrast between “two cultures” of statistical modeling, which he labels the “data modeling culture” and the “algorithmic modeling culture”. We identify the latter with what is commonly referred to as AI, although we recognize that there is no widely accepted division between what is part of statistics and what is part of AI. Breiman outlines what he sees are the broad differences between the two approaches, but a key feature for our purposes is that the former, typically informed explicitly by scientific theory, uses a parametric data model. This model will typically specify the dependent variable, the predictor variables, and the functional form of the relationship between the dependent and the predictor variables and the stochastic form of the disturbance term in the model.

We now consider the effect of having access to AI as an alternative tool for statistical modeling of the landscape. Breiman (2001, p. 205) describes the alternative “algorithmic” approach as follows:

The approach is that nature produces data in a black box where the insides are complex, mysterious and, at least, partly unknowable. What is observed is a set of \mathbf{x} 's and a subsequent set of \mathbf{y} 's that come out. The problem is to find an algorithm $f(\mathbf{x})$ such that for a future \mathbf{x} in a test set, $f(\mathbf{x})$ will be a good predictor of \mathbf{y} .

The data-generating processes in many combinatorial-type problems – drug discovery, materials science, genomics, etc. – do appear to fit the description of “complex, mysterious and, at least, partly unknowable.” To the extent that the AI approach provides a better prediction model (in at least some circumstances) for valuable new combinations, it has the potential to boost the productivity of the innovation search process, causing the probability ranking curve to swivel in a clockwise direction as in Figure 2, indicating a more discriminating prediction model. It is important to note that textbooks on AI (machine learning) typically subsume parametric regression and classification models as forms of machine learning. However, we find Breiman’s distinction to be useful in highlighting the shock to the innovation process that the rapid advance in AI (specifically, the subfield of machine learning) has engendered, and take his narrower category to be what we refer to as AI (Athey and Imbens, 2019).

How does AI fit our generic predict-test workflow? We assume that part of the knowledge base consists of data on previous experiments (or market tests) that we treat for simplicity as indicating tested combinations were successes or failures. These binary outcome data together with input data on the combinations (amino acid sequences, molecular descriptors, etc.) are the training data for our AI. The measure of a good prediction model will be how well it predicts (or generalizes) outside of the training sample.

Of course, there is a vast – and rapidly growing – array of available AI algorithms. To give a flavor of their use in generating prediction models we note just a selection here and relate them to our generic workflow of classifying potential combinations into predicted successes and failures. Probably the most intuitive algorithm is k-nearest neighbors. The predicted probability of success of a candidate combination is simply the average success rate of the k nearest neighbors in the search

space. Decision tree methods work by instead segmenting the search space into success and failure regions. The more complex decision tree methods used in practice (e.g., random forests) involve multiple trees that are combined together to produce the predicted probability of success. A third example is the naïve Bayes classifier. In contrast to the approaches noted already, this approach concentrates on the individual “features” – the states of an element of the string describing a combination in our generic example – which are each assumed to have an independent effect on the success of a combination. The output is an estimate of the probability of success conditional on the states of all the elements in the string describing that candidate combination.

In addition to the increased availability of data, much of the recent excitement concerning AI is due to a rapid improvement in algorithms. Of particular note is the rapid improvement in prediction models based on artificial neural networks, most notably so-called deep learning algorithms.

Deep learning is making major advances in solving problems that have resisted the best attempts of the artificial intelligence community for years. It has turned out to be very good at discovering intricate structure in high-dimensional data and is therefore applicable to many domains of science, business, and government. (LeCun et al., 2015, p. 436)

Although an in-depth discussion of the technical advances underlying deep learning is beyond the scope of this paper, three aspects are worth highlighting. First, the development and optimization of multilayer neural networks allows for substantial improvement in the ability to predict outcomes in high-dimensional spaces with complex non-linear interactions (LeCun et al., 2015).⁹ Second, given that previous generations of machine learning were constrained by the need to extract features (or explanatory variables) by hand before statistical analysis, a major advance in machine learning involves the use of ‘representation learning’ to automatically extract the relevant features.¹⁰ And

⁹For example, a recent review of the use of deep learning in computational biology notes that the “rapid increase in biological data dimensions and acquisition rates is challenging conventional analysis strategies,” and that “[m]odern machine learning methods, such as deep learning, promise to leverage very large data sets for finding hidden structure within them, and for making accurate predictions” (Angermueller et al., 2016, p. 1). Another review of the use of deep learning in computational chemistry highlights how deep learning has a “ubiquity and broad applicability to a wide range of challenges in the field, including quantitative activity relationship, virtual screening, protein structure prediction, quantum chemistry, materials design and property prediction” (Goh et al., 2017).

¹⁰As described by LeCun et al. (2015, p. 436), “[c]onventional machine-learning techniques were limited in their

third, recent optimism about developments in deep learning relates to demonstrated out-of-sample performance of deep learning models across a range of tasks, including image recognition, speech recognition, language processing, and autonomous vehicles.¹¹

6 A Multi-Stage Discovery Process and the Option to Abandon

In Section 4, we assumed a highly simplified two-stage innovation process involving prediction (development of a prediction model choosing a testing threshold) and testing (testing all combinations with a probability of success at or above a threshold value that depended on the market value of an innovation success and the cost of conducting the test). The motivation for this simple set-up is that the discovery pipeline often involves the production of a priority list for testing where that testing is expensive. We introduced AI in Section 5 as a possible way to improve the list.

In reality, innovation is more complex and involves multiple stages rather than just two. The typical tasks (or stages) of the drug discovery process include target identification, hit generation, generation of lead compounds from the hits (“hits to lead”), optimization of the lead compounds, pre-clinical trials (animal studies), and phase I, II, and III human clinical trials. In materials discovery, the stages can involve prediction of molecule properties, synthesis of the molecules and characterization of the actual properties through testing, and this can be followed by further stages such as investigating the ability to synthesize at scale and testing the molecule under the different environmental conditions that could be observed in the field.

In this section, we extend our two-stage prediction-test model to a multi-stage model. The discovery process now involves S stages: the initial prediction stage, the final test stage and $\tilde{S} = S - 2$ intermediate stages. The intermediate stages can be thought of as involving some combination of prediction refinements and preliminary tests. Our search process is an example of what Roberts

ability to process natural data in their raw form. For decades, constructing a pattern-recognition or machine-learning system required careful engineering and considerable domain expertise to design a feature extractor that transformed the raw data (such as the pixel values of an image) into a suitable internal representation or feature vector from which the learning subsystem, often a classifier, could detect or classify patterns in the input. . . . Representation learning is a set of methods that allows a machine to be fed with raw data and to automatically discover the representations needed for detection or classification.”

¹¹While scholars consider deep learning models something of a black box due to their complexity, recent theoretical work has made progress in understanding why these models have had such success in generalizing beyond their training data (see, for example, Shwartz-Ziv and Tishby (2017)).

and Weitzman (1981) call a *Sequential Development Project*: costs are additive across stages, value is received only at the end of the project, and there is an option of abandoning the project at the end of each stage. Our main focus is on the implications of the option to abandon on the robustness of the result from the two-stage model that an improvement in the prediction model (as captured by an increase in b) will increase the expected total net value of innovation.

A tractable assumption for a multi-stage discovery process is for the expected gross value of any given combination to evolve as a simple discrete-stage random-growth Markov process with equal probabilities of an up and down move at any intermediate stage. More specifically, for each intermediate stage, we assume the outcome of that stage follows a zero mean random growth process for which there is a 0.5 probability of a proportionate increase in the estimated probability of success equal to u and a 0.5 probability of a proportionate decrease in the estimated probability of success also equal to u .¹²

To minimize notation we assume, without loss of generality, that the gross payoff from a success, π , is equal to 1. Thus the expected gross value of a combination at the completion of any given stage is just the estimated probability of success at that stage. We assume all stages must be completed to launch an innovation on the market so that it is not possible to skip a stage. At the completion of the final (test) stage, the gross value is equal to 1 (a success) or 0 (a failure) with the estimated probability of a success going into the testing stage equal to the estimated probability of success at the completion of the last intermediate stage.

With this random process, the possible paths of the gross value of a combination gives rise to a binominal lattice decision tree structure. An example of this structure for a five-stage discovery process is shown in Figure 6. The tree has the convenient property that the branches recombine, producing the lattice. Given this lattice (or recombinant) structure, the number of possible outcomes (nodes) is equal to the number of stages, so that the number of nodes only grows linearly as the stages advance, greatly reducing the computational burden. For example, at Stage 4, there are four nodes and thus four possible values for the estimated probability of success: $(1+u)^3p_r$, $(1+u)^2(1-u)p_r$, $(1+u)(1-u)^2p_r$, and $(1-u)^3p_r$. However, there can be multiple paths

¹²We assume we are always operating in a range of initial estimated probabilities such that the estimated probability of success at the completion of any stage never exceeds 1, which requires that $u \leq \sqrt[5]{p_r^{-1} - 1}$.

to a given node, where the number of paths to a given node is given by the binomial coefficient:

$$\binom{s-1}{n_s-1}, \quad (12)$$

where n_s is the number of nodes (ordered from the top in the lattice) at stage s . For example, at Stage 4, there is one path to $(1+u)^3$, three paths to $(1+u)^2(1-u)$, three paths to $(1+u)(1-u)^2$ and one path to $(1-u)^3$. When viewed from the end of the first (prediction) stage for a combination with an estimated probability of success equal to p_r a useful implication of the binomial lattice structure is that the expected probability of success (and thus expected gross value) at the completion of any subsequent stage remains at p_r . For stage s (for $s = 2, \dots, S-1$), this expected probability of success can be calculated as:

$$p_r 0.5^{s-1} \sum_{n_s=1}^s \left(\binom{s-1}{n_s-1} (1+u)^{s-1-(n_s-1)} (1-u)^{n_s-1} \right) = p_r. \quad (13)$$

A further noteworthy feature of the binominal lattice structure is also evident from Figure 6: the uncertainty as to the revealed value of a given combination when viewed from the end of the prediction stage (say as measured by the possible range or variance of the outcomes) is increasing in the number of intermediate stages. For example, in a five-stage process, the range of possible values increases from $[(1+u)p^r, (1-u)p^r]$ to $[(1+u)^3p^r, (1-u)^3p^r]$ as we compare Stage 2 and Stage 4.

As noted, a key feature of our multi-stage model is that the innovator has the option to abandon a combination after any stage. The cost of completing a stage after the first prediction stage is c_s . The cost of conducting a given stage is assumed to be the same for all combinations. Once a stage is completed the cost of that stage is sunk. The total cost of any combination that successfully makes it all the way through the discovery pipeline is:

$$\sum_{s=2}^S c_s. \quad (14)$$

Given these assumptions on the evolution of value and cost as a combination advances through the pipeline, the innovator can determine whether to advance a combination to the next stage of the

pipeline by backward induction starting from the final decision stage – i.e., the decision whether or not to send the combination for final testing. The optimal decision at any node that the innovator might reach is based on the assumption that the innovator will act optimally at any subsequent node. In particular, the optimal action at any node will take into account the option to abandon that may later be exercised if there are sufficiently unfavorable updates to the estimated probability of success at later stages.

We illustrate the decision making process in a multi-stage discovery process with a three-stage process – i.e. a process with just a single intermediate stage. The logic of the three-stage process can be extended to any S -stage process. Given the ranking function from the first prediction stage, we focus on the decision as to which combinations to advance to the intermediate stage. We assume (as in Case 1 in Section 4) that the objective of our expected value maximizing innovator is to find all combinations with an expected net value greater than or equal to zero. Of particular interest will be how the decision to advance combinations along the pipeline is affected by uncertainty (which depends on u) and costs at the intermediate and final stages (c_2 and c_3 respectively). We examine how the option to abandon affects both the expected number of combinations that will be tested and the expected net total value of discovery for any given post-prediction (logistic) ranking function.

As a starting point, it is useful to consider the decision of innovator if there is no option to abandon. The marginal expected gross value of a combination is simply:

$$0.5(1 + u)p_r + 0.5(1 - u)p_r = p_r, \tag{15}$$

where we assume that $(1 + u)p_r \leq 1$. The expected value maximizing innovator will again choose how far down the ranking function to move by choosing to advance all combinations for which the expected marginal gross value is greater than the marginal cost, where the latter is now $c = c_2 + c_3$. The probability threshold is therefore now equal to $c_2 + c_3$. This is shown in Figure 7. Thus the only difference from the case examined in Section 4 is that marginal cost is higher (the cost of the intermediate stage plus the cost of the testing stage), which means fewer combinations might advance from Stage 1 given the additional cost involved in having to conduct the intermediate

stage.

We next introduce an option to abandon if there is a bad realization at the intermediate stage (i.e. a realization leads to a revised estimated probability of success equal to $(1 - u)p_r$). Given that the Stage 2 cost is now sunk, the option to abandon will be exercised provided:

$$\begin{aligned} (1 - u)p_r &< c_3 \\ \Rightarrow p_r &< \frac{c_3}{1 - u}. \end{aligned} \tag{16}$$

As shown in Figure X, there can be combinations that would have advanced to testing without the option will now be abandoned after the intermediate stage provided:

$$\begin{aligned} \frac{c_3}{1 - u} &> c_2 + c_3 \\ \Rightarrow u &> \frac{c_2}{c_2 + c_3}. \end{aligned} \tag{17}$$

The left-hand-side of the last expression can usefully be viewed as the proportionate reduction in the marginal expected value in the bad realization. The right-hand-side can be viewed as the proportionate reduction in relevant costs due to the fact that Stage 2 costs are now sunk. This reduction in relevant costs could lead to projects continuing that would not have gone to the intermediate stage if the bad realization on the estimated probability of success had been known in advance. However, the condition shows that some projects that would move to the intermediate stage will be abandoned assuming a high enough value of u . In Figure 6, we assume that the condition in equation (4) holds, which can alternatively be viewed as the necessary condition for the option to abandon actual having (option) value for the innovator.

For combinations that meet the condition in equation (17), the innovator will thus assume that they would be abandoned after the intermediate given a bad realization on the estimated probability of success. Viewed from Stage 1, for these combinations the expected marginal gross value is therefore: $0.5(1 + u)p_r$. And the *expected* marginal cost is: $c_2 + 0.5c_3$.

There is thus a discrete downward jump in both the expected gross marginal value and expected marginal cost curves when we enter into the probability range where the option to abandon would be exercised when the bad realization occurs. The position of these curves will again determine which (and how many) combinations advance to the intermediate stage. The condition required for a combination to enter the intermediate stage is:

$$\begin{aligned} 0.5(1+u)p_r &\geq c_2 + 0.5c_3 \\ \Rightarrow p_r &\geq \frac{2c_2 + c_3}{1+u}. \end{aligned} \tag{18}$$

Comparing equation (16) and equation (18), it is easy to verify that the condition for the probability threshold to be lower when the option to abandon is available (and thus the condition for more combinations to advance to the intermediate stage) is given by equation (17), which we have assumed to hold.

Substituting equation (18) into the ranking function (2) and again assuming for simplicity that equation (18) holds with equality, we obtain the optimal number of combinations to advance to the intermediate stage:

$$r^{**} = G - \frac{\ln\left(\frac{2^A - D - G}{G}\right) - \ln\left(\frac{1+u}{2c_2+c_3} - 1\right)}{b}. \tag{19}$$

To the extent that the option to abandon has value, the expected total net value to the innovator will be higher than when no such option exists when viewed from the end of the first (prediction) stage. An increase in uncertainty will raise the number of combinations that advance to the next stage of the pipeline and also raise this expected value.

Moreover, an increase in either intermediate or final stage costs will lower the number combinations that advance, with the sensitivity to the intermediate-stage cost being greater than to the final-stage cost. This difference in sensitivity is explained by the fact that there is only a 0.5 chance that the marginal combination will advance to the final stage given that the option to abandon will be exercised in the event of a bad realization, which has a 0.5 chance of occurring.

We finally consider how an improvement in the prediction model (again captured by an increase in b) will affect both the number of combinations advancing through the pipeline and also expected total net value, V^e . The effect of such an improvement is most easily seen from Figure 7. An increase in b will shift the expected marginal gross value curve upwards. For a large enough shift, it is apparent that the number of combinations that advance will increase and also that V^e will increase. As will be familiar from Section 4, the increase in expected value is due to the increase in net value on the intra-marginal combinations (evaluated at r^*) and any additional increase due to an increase in the optimal number of combinations that advance from r^* to r^{**} .

The logic of our three-stage process can be extended to an S -stage process. Entering the penultimate stage (i.e. the last intermediate stage) the analysis is identical at any given node to the three-stage example given above. Any positive option value for this sub-problem will result in positive option value for the decision problem overall. A similar argument applies to any other node along the decision tree, where the decision maker assumes they will act optimally in terms of any decision to abandon at any subsequent node. Therefore, if the option to abandon has value at any node in the decision tree, then the existence of the option has value for the decision process overall.

Although our analysis has been conducted in the context of a specific tractable example, it hopefully illustrates how an AI-induced improvement in the prediction model can increase the productivity of a multi-stage discovery process even where the use of AI only improves the performance at the initial prediction stage and combinations can be abandoned at any stage as new information becomes available.

In the next section, we discuss some more general issues that take particular salience in the context of a multi-stage process: bottlenecks, autonomous discovery systems, and the demand for specific (possibly stage-related) R&D skills.

7 Bottlenecks, Autonomous Discovery Systems, and Demands for R&D Labor

7.1 Bottlenecks

In both our two-stage and multi-stage discovery processes, an improvement in the prediction model leads to an increase in the expected total net value of innovation. However, the size of this positive effect depends on the costliness of conducting later stages of the process.

The relevance of later stage costs are most easily seen in our two-stage model. In Case 1 (search for all combinations with an expected net value greater than or equal to zero), a higher cost of testing unambiguously reduces the increase in expected total net value, with the impact on the intra-marginal combinations – i.e., those combinations that would go for testing even without the improvement in the prediction model – likely to be of particular benefit (see Figure 3).

The situation is more ambiguous in Case 2 (search for a single successful combination). There the impact of a lower cost of testing on the expected total net value depends on whether the marginal cost curve intersects above the crossing point of the expected marginal gross value curves (Figure 4b) or below that crossing point (Figure 4c). In the former case, a decrease in cost raises the net benefit obtained from the improved prediction model; in the latter it lowers it.

The difference is even more pronounced if there are actual binding capacity constraints on later stages that prevent the innovator from taking advantage of an improved prediction model to increase the number of combinations that move down the discovery pipeline. For example, in the case of drug discovery there may be a capacity constraint on the number of clinical trials that can be run or in the case of materials discovery a limited availability of the necessary equipment to conduct experimental tests on novel molecules. Thus, it is possible that “bottlenecks” – thought of as either high later-stage costs or binding capacity constraints – will attenuate the benefits of the improved prediction model.

As they relate to the innovation process, one characterization of recent advances in machine prediction is as a new general purpose technology (GPT) for invention (Agrawal et al., 2019a; Cockburn et al., 2019). As has been identified in other contexts, new GPTs may only have their

full effect after a significant elapse of time due to the need for complementary upstream and downstream investments. In a classic paper, Paul David (1990) explored the lagged productivity effects of both the invention of distributed electricity and the computer. Drawing on historical analogies of delayed productivity effects, David cautioned against undue pessimism due to the apparent limited impact of computers on productivity, notwithstanding their growing prevalence: “Closer study of some economic history of technology, and familiarity with the story of the dynamo revolution in particular, should help us avoid the pitfall of undue sanguinity and the pitfall of unrealistic impatience into which current discussions of the productivity paradox seem to plunge all too frequently” (p. 359-360).

In the example of drug discovery, the benefit of AI in generating and ranking a large number of leads may be compromised, say, by a capacity constraint at the following lead optimization stage. It is worth considering a David-type solution that might evolve over time. As machine prediction develops, it might also be applied directly in other stages, including substituting for human judgment in lead optimization (see Agrawal et al., 2019b, on the roles of prediction and judgement in a multi-stage decision process).

Much of the recent interest in AI in the econometrics and policy evaluation literature has focused on the value of machine learning in causal inference and policy evaluation (see Athey (2019), for a recent survey). Such tools may be especially relevant in later stages of a multi-stage discovery process as complements, say, to random and natural experiments. Thus AI may itself help relieve the bottleneck problem over time.

Machine learning is being applied to help control for confounding variables to estimate average treatment effects and also to estimate heterogeneous treatment effects (Athey, 2019; Athey and Imbens, 2019). Kleinberg et al. (2015) highlight a very direct way that improved prediction can improve later stage evaluation. Taking the example of joint replacement, they observe that the expected net benefit of difficult but costly and painful surgery to the patient depends on how long they can expect to live:

The payoff to surgery depends on (eventual) mortality, creating a pure prediction problem. Put differently, the policy challenge is: can we predict which surgeries will be

futile using only data available at the time of surgery? This would allow us to save both dollars and disutility for patients. (Kleinberg et al., 2015, p. 493).

7.2 Autonomous discovery systems

The slow speed and high failure rates of discovery pipelines plagued with bottlenecks has increased interest in more autonomous processes that utilize a combination of AI and robotics. For example, AI-based predictions determine which candidates are tested using robotic high throughput screening (HTS) methods. An early example of such an autonomous system is the Robot Scientist (Sparkes et al., 2010). The first prototype Robot Scientist, Adam, generated hypotheses and carried out experiments related to the functional genomics of a yeast. While it is unlikely that closed-loop systems will lead to a mass replacement of scientists, AI and automation may ease bottlenecks across the discovery pipeline.

Lamenting the slow speed and high cost of the development and deployment of advanced materials using the traditional approach – new materials typically reach the market after 10-20 years of basic and applied research – Tabor et al. (2018, p. 15) outline what they see as required for an autonomous (closed-loop) innovation process:

To fully exploit the advances in robotics, machine learning, high-throughput screening, combinatorial methods, and in situ or in operando characterization, we must close the loop in the research process. This means that humans must partner with autonomous research robots to design experimental campaigns and that the research robots perform experiments, analyze the results, update our understanding and then use AI and machine learning to design new experiments optimized to the research goals, thus completing one experimental loop.

Current processes are affected by the bottleneck problem, which Tabor et al. (2018, p. 16) suggest is in the “experimental synthesis, characterization and testing of theoretically proposed materials,” but that if an autonomous approach could be implemented “the bottleneck would move to AI.”

The prediction model in our two-task process played a basic role: predicting the probability of success of any given candidate combination and thus informing which combinations should go for testing. In a truly closed-loop process, an additional source of value from tests needs to be considered: how the data from those tests can be used to improve the prediction model. For example, an additional test observation that allows for more accurate discrimination for candidate combinations in the neighborhood of the probability threshold may have a particularly high information value as the innovator seeks to exploit promising parts of the combinatorial space.

In contrast, test observations whose neighbors are in parts of the space that have very high or very low probabilities of success may have low information value. On the other hand, when the cost of autonomous experimental tests is low, it may make sense to conduct more exploratory testing of regions of the space where data on successes and failures is sparse, potentially opening up truly novel approaches (say in materials for energy harvesting or storage). While

the best experimental design will be context dependent, the important point is that a truly autonomous process will require closing the loop by choosing which experiments to conduct based on both the direct expected *innovation* value and the indirect expected data value of the test.

7.3 R&D labor demand effects

Much of the economic analysis of AI has focused on its potential implications for labor demand and consequently for wages, labor shares, and employment. Concerns of adverse effects on scientists, engineers, and others working in innovation intensive sectors will be heightened if we move towards truly autonomous discovery systems. We therefore finally consider the implications of the increased use of AI in discovery on the demands for various types of R&D labor.

In a series of papers, Daron Acemoglu and Pascual Restrepo extensively analyzed the effects of AI and robotics on labor demand (see, e.g., Acemoglu and Restrepo, 2019a, 2018). We adopt the useful framework set out in Acemoglu and Restrepo (2019b). This framework decomposes the labor demand effects into displacement effects, countervailing productivity effects, and new task (or reinstatement) effects.

Our multi-stage model clearly shows the potential for human skill *displacement* as a result of

breakthroughs in AI in both of the tasks. As outlined in Section 5, prediction can be achieved through various methods including theory, computer-based simulation, and educated guesses. Indeed, these forms of prediction can be present together as, say, intuition guides theory development and predictions are generated from complex theoretical models using computer aided simulations as, for example, in the use of density functional theory in chemistry. Statistical data modeling can also be the basis for prediction by highlighting associations in the data and can also inform theory and simulation (e.g. by aiding the initial calibration of simulation models).

By providing an alternative source of predictions, there is an obvious potential for AI to displace certain existing skills. Our multi-stage model also points to the potential for displacement in the testing stage.

However, while we might normally expect that improved prediction would increase the demand for testing through an increased flow of promising combinations through the discovery pipeline, Case 2, search for a single successful combination, of the two-stage model of Section 4, shows that the effect of an upward shift in the ranking function on the amount of testing depends on the cost of testing relative to the value of a successful innovation. At a sufficiently low levels this ratio, a more discriminating prediction model actually leads to fewer tests performed as the expected value of the marginal test declines. The reason is that improved prediction increases the probability that the target will have achieved by the other tests performed, lowering the expected value of the marginal test.

While the potential for displacement is clearly present, there may also be *countervailing productivity effects* that increase the demand for certain skills. Most obviously, where improved prediction does increase the flow of combinations through the pipeline, the demand for the skills of testers (say experts in experimental design or lab technicians) will increase. Acemoglu and Restrepo (2019b) also note additional sources of countervailing effect through increased capital investments and the “deepening” of automation. These processes may be particularly relevant as investments and technological development takes place to remove the bottlenecks in the system in order to better take advantage of the new prediction GPT.

Finally, Acemoglu and Restrepo (2019b) stress the importance of the demand for new tasks that

come into being as a result of AI. The AlphaFold example with which we open the paper provides a good example: this protein folding prediction effort brought together an interdisciplinary team from structural biology, physics, and machine learning. Critical to the success of teams at the cutting edge of scientific discovery and innovation is likely to be the ability of members to communicate with other specialists: AI specialists need to have a sufficiently deep understanding of biology to work with the structural biologists; the structural biologists will need to know enough machine learning for tasks such as preparing the data for building well-performing prediction models to complement the AI specialists. Skill sets that combine expertise such as computer programming or statistics with knowledge of specific scientific domains are likely to be in especially heavy demand. Such evolving demands will lead to incentives for specific upskilling and career transitions and also changes in the curricula (and hiring priorities) in universities and other training programs. We return to the policy and management implications of an increased use of AI in the discovery process in the concluding section.

8 Spillovers: Ideas versus Data

In this section, we briefly return to the spillover mechanism that has been at the centre of much of modern growth theory. As outlined in Section 3, a key mechanism sustaining economic growth in many endogenous growth models is that the discovery of new ideas that spill over to expand the feedstock of existing ideas and support the discovery of yet further ideas. The Cobb-Douglas form has been widely used as a functional form capturing the relationship between the stock of ideas and innovation (see, e.g., Jones, 1995, 2005):

$$I = \dot{A} = \beta A^\phi, \quad \text{where } \beta > 0 \text{ and } 0 < \phi \leq 1. \quad (20)$$

Romer's original model (Romer, 1990) assumed $\phi = 1$. With $\phi < 1$ (as in Jones, 1995), growth will converge to zero (in the absence of population growth) as the stock of ideas grows to infinity.

The combinatorial view makes explicit the mechanism by which existing ideas are drawn on to produce new ideas: each idea is a combination of other ideas and increases in the stock of

existing ideas causes the set of potential combinations, 2^A , to grow exponentially. In contrast to the declining growth rate in the Jones model, the exponential growth in potential combinations causes an almost embarrassment of riches that is poorly matched to the recent disappointing growth record (see Bloom et al., 2017).¹³

We thus focus on a second possible spillover mechanism that is suggested by the combinatorial model. To isolate this mechanism, we make the extreme assumption that the fruitfulness of the expanded space in terms of valuable innovations as A increases is zero in the vicinity of the current idea stock. However, the discovery process of each successful idea is further assumed to be associated with on average α failures. More specifically, the innovation search process produces data according to:

$$D = (1 + \alpha)A. \tag{21}$$

We assume that the performance of the innovation search process improves as increased data availability on successes and failures allows the production of improved prediction models. Thus, although the fruitfulness of any new parts of the combinatorial space is assumed to be zero as the stock of ideas grow, an improved search process increases the chances of finding successful combinations in the older (fruitful) space is increased. Again assuming a Cobb-Douglas functional form with $\phi < 1$, the innovation function is assumed to take the simple form:

$$I = \dot{A} = D^\phi = (1 + \alpha)^\phi A^\phi = \beta A^\phi \quad \text{where } \beta = (1 + \alpha)^\phi. \tag{22}$$

Innovation again increases with the stock of ideas, but the spillover mechanism is through the value of increased data availability for producing improved prediction models. The assumption that $\phi < 1$ reflects diminishing marginal returns to data (see Varian, 2019). Figure 8a shows the relationship between data and innovation performance (reflecting the positive but diminishing marginal impact of data on the quality of available maps); Figure 8b shows the corresponding implied relationship between the stock of ideas (given equations (21) and (22)) and innovation

¹³See Agrawal et al. (2019a) for an examination of functional forms for the innovation function in a combinatorial growth model.

performance. A once-off improvement in prediction models (say because of a major improvement in a machine learning algorithm) will increase the value of ϕ and cause an upward shift in both curves and also an increase the growth rate along the transition path to the steady state. The ImageNet competition - where the error rate in image recognition dropped dramatically due to improved algorithms notwithstanding no change in the underlying data - provides an illustration of the impact of improved prediction models on prediction task performance (Varian, 2019). Moreover, to the extent that access to data on failures (α) is a variable, improved access to this form of data will raise the value of β and increase the growth rate along the transition path. Improved prediction algorithms and improved access to data on failures will interact positively in terms of their impact on the growth rate. We elaborate on the possible policy implications of the data-related spillover mechanism in the concluding comments.

9 Concluding Comments

AI has achieved well-documented recent successes in tasks such as image recognition, speech recognition, language translation, recommendation systems and autonomous vehicles. It is now viewed as a GPT for prediction tasks that can be combined with other components to produce novel products and services (Agrawal et al., 2019a). AI-based prediction models can deal with vast combinatorial spaces (e.g., pixels in an image). Our focus in this paper is on AI in the discovery process itself. Conceptualization of the innovation process as search over a combinatorial space highlights the value of AI in helping to map this space in the form of a prediction model.

The burgeoning scientific literature applying AI tools suggests the practical value of being able to make predictions when the number of potential combinations is in the billions. But prediction is just one part of the discovery process. We thus embedded the prediction task as just one part of a multi-stage process. In our model of AI-aided discovery the prediction model produced a ranking function that essentially provided a priority list for later (expensive) testing. Improvements in the prediction model – say as a result of the availability of a better performing algorithm – allowed for a more discriminating prioritization and ultimately for a more productive discovery pipeline.

The main testable hypothesis from the model is therefore that access to AI will increase the

productivity of the innovation process for combinatorial-type problems. Although there has been rapid growth in the use of AI – and a great deal of optimism expressed about its effects – the productivity benefits have yet to be well established. Indeed, there is skepticism about its ultimate benefits, in part reflecting the previous waves of optimism and pessimism associated with AI. There have also been well-publicized cases where early optimistic scientific findings have been reversed and exits of high profile research groups. Skeptics have a number of concerns, including: error-prone data leading to poorly performing models; failure to replicate results; excessive concentration on already well explored regions of the search space where data are plentiful; difficulties of interpretation of “black-box” models and associated issues of trust; and models that appear to work well on test data sets but ultimately perform poorly due to redundancy between training and test data (Wallach and Heifets, 2018).

The next step in our research program is therefore to look for evidence of the productivity effects of increased access to AI-aided discovery. This is obviously complicated by the relative newness of some of the main developments (including the improved performance of deep neural network algorithms). However, we hope it will be possible to exploit plausibly exogenous variation in the access to AI tools. One idea we are exploring is to use variation in historic expertise at the department level in machine learning in computer science and engineering departments, and to see how it relates to later output in application domains within the university (medicine, chemistry, etc.) based on assumption of local knowledge spillovers. Another avenue is to look for shocks to access – say the availability of the AtomWise AIMS service discussed in the introduction – and see how it affects the productivity of researchers in a difference-in-difference framework.

The rise in AI-aided scientific discovery and innovation also raises some novel policy and managerial challenges. Looking through the lens of the model, we finally reflect briefly on some of these challenges as they affect academia, government and industry.

Perhaps the most consequential implication of our combinatorial search model is the importance of data. Access to data – findings on prior successes and failures in our simplified setting – is essential to building effective prediction models. The paper also highlights the importance of data spillovers to the process of sustaining economic growth. Greater access to data allows better exploitation of

the fruitful parts of the space by building better maps of that space.

Scientists are increasingly recognizing the importance of open data, and numerous collective efforts have arisen to improve data access. Funding agencies are also more often requiring open access as a condition for funding. The number of open source datasets is growing. Examples include *The 1000 Genomes Project* (genomics), *PubChem* (chemistry) and *The Materials Project* (materials science).

AI is itself being applied to help discover data. For example, *BenchSci* is an AI-based search technology for identifying antibodies that act as reagents in scientific experiments. It attempts to deal with the challenge that reports on compound efficiency are scattered through millions of scientific papers with limited standardization on how the reports are provided. *BenchSci* extracts compound efficiency information thus allowing scientists to more effectively identify compounds for experiments. Another example is *Meta*^α, an AI-based search technology for identifying relevant scientific papers and tracking the evolution of scientific ideas.

A significant concern is a lack of access of data on failures, which are more likely to remain hidden in the notebooks of experimentalists given the much greater likelihood of publishing data on successes (See, e.g., Raccuglia et al., 2016). The successful training of supervised machine learning models requires knowledge of what failed as well as what succeeded in the past. Another concern stems from the importance of private companies in the innovation process. Private companies may have strong incentives to limit access to proprietary data. This gives larger companies an advantage that can become self-reinforcing as they monopolize access to expanding datasets. However, many companies are recognizing the advantage of open innovation models (See Chesbrough, 2006, for an early discussion) as they gain access as well as provide access by engaging in open knowledge sharing networks. Firms may also have an incentive to make data (and algorithms) accessible as the attempt to stave off regulation and anti-trust enforcement. Furthermore, the increasing salience of privacy issues could also limit access to certain forms of data such as patient health records, although privacy issues should not be a constraint where personal data is not involved (e.g. datasets on molecular descriptors and properties).

The model also highlights the problem of bottlenecks in the discovery process that could limit

the productivity benefits of machine learning as a GPT for prediction. The issue of bottlenecks has probably received most attention in the context of materials discovery. The long timelines between the initiation of research and the launch of new materials has led to interest in more autonomous (or “self-driving”) discovery processes (Aspuru-Guzik and Persson, 2018). The promise is expressed by Tabor et al. (2018, p.5): “[P]latforms that integrate AI with automated and robotized synthesis and characterization have the potential to accelerate the entire materials discovery and innovation process to reduce this time [10-20 years] by an order of magnitude.” To deliver on this promise, they call on efforts and leadership from academia, governments and industry in building the required multi-disciplinary workforce and supporting collective action efforts such as the Material Genome Initiative and the multi-country Mission Innovation Initiative.

Of course, the move to such autonomous systems involving an increased use of AI and robotics raises the specter of significant displacement of parts of the R&D workforce. As emphasized by Acemoglu and Restrepo (2019), the key countervailing force for this displacement effect is likely to be the demand for new skills. These new skills are likely to be cross-cutting – for example, biologists that know machine learning and machine-learning experts that know biology. A key skill will be the ability to work in multi-disciplinary teams. This creates challenges for academia, government and industry as they engage in workforce planning (see, for example, the Future Jobs Ireland initiative and Canada’s Future Skills Council as responses). Such integration challenges have long been recognized in innovation-intensive industries such as biotechnology.

To perform well the [biotechnology] sector requires appropriate mechanisms for bringing together and integrating the right mix of cross-disciplinary talents, skills and capabilities. These mechanisms include organizational structures and strategies and the means by which different types of organizations (large firms, small start-ups, universities, etc.) interact. However, the sector also requires micro-organizational mechanisms for creating truly integrated problem solving and avoiding islands of specialization. And, perhaps most importantly, requires ways of getting together the right mix of people from different scientific and functional backgrounds to collaborate and exchange information. (Pisano, 2006, p. 76)

For universities the challenge is to allocate resources to ensure that they are at the forefront of these emerging developments in the discovery process. This means research efforts to ensure the flow of new ideas that take advantage of (or indeed build) the new GPTs for discovery and also to ensure appropriate research-led teaching by frontier researchers so that the innovation workers of the future have the skills that will be newly in demand.

References

- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, tasks and technologies: Implications for employment and earnings,” in *Handbook of labor economics*, Elsevier, vol. 4, 1043–1171.
- ACEMOGLU, D. AND P. RESTREPO (2018): “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108, 1488–1542.
- (2019a): “Artificial Intelligence, Automation, and Work,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.
- (2019b): “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33, 3–30.
- AGHION, P. AND P. HOWITT (1992): “A Model of Growth through Creative Destruction,” *Econometrica*, 60.
- AGHION, P., B. JONES, AND C. JONES (2019): “Artificial Intelligence and Economic Growth,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.
- AGRAWAL, A., A. GOLDFARB, AND F. TEODORIDIS (2016): “Understanding the changing structure of scientific inquiry,” *American Economic Journal: Applied Economics*, 8, 100–128.
- AGRAWAL, A., J. MCHALE, AND A. OETTL (2019a): “Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.
- AGRAWAL, A. K., J. S. GANS, AND A. GOLDFARB (2019b): “Exploring the impact of artificial intelligence: Prediction versus judgment,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.
- ANGERMUELLER, C., T. PÄRNAMAA, L. PARTS, AND O. STEGLE (2016): “Deep learning for computational biology,” *Molecular systems biology*, 12, 878.
- ARROW, K. J. (1962): “The economic implications of learning by doing,” *The review of economic studies*, 29, 155–173.
- ARTHUR, B. W. (2009): *The Nature of Technology: What it is and How it Evolves*, Penguin Books, London.
- ASPURU-GUZZIK, A. AND K. PERSSON (2018): “Materials Acceleration Platform: Accelerating Advanced Energy Materials Discovery by Integrating High-Throughput Methods and Artificial Intelligence.” Tech. rep., Canadian Institute for Advanced Research.
- ATHEY, S. (2017): “Beyond prediction: Using big data for policy problems,” *Science*, 355, 483–485.
- (2019): “The impact of machine learning on economics,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.

- ATHEY, S. AND G. IMBENS (2019): “Machine Learning Methods Economists Should Know About,” *arXiv preprint arXiv:1903.10075*.
- AUTOR, D. H. (2015): “Why are there still so many jobs? The history and future of workplace automation,” *Journal of Economic Perspectives*, 29, 3–30.
- BLOOM, N., C. I. JONES, J. VAN REENEN, AND M. WEBB (2017): “Are Ideas Getting Harder to Find?” Working Paper 23782, National Bureau of Economic Research.
- BREIMAN, L. (2001): “Statistical Modeling: The Two Cultures,” *Statistical Science*, 16, 199–231.
- BUTLER, K. T., D. W. DAVIES, H. CARTWRIGHT, O. ISAYEV, AND A. WALSH (2018): “Machine learning for molecular and materials science,” *Nature*, 559, 547.
- CHADE, H. AND L. SMITH (2006): “Simultaneous search,” *Econometrica*, 74, 1293–1307.
- CHEN, H., O. ENKVIST, Y. WANG, M. OLIVECRONA, AND T. BLASCHKE (2018): “The rise of deep learning in drug discovery,” *Drug discovery today*.
- CHESBROUGH, H. W. (2006): *Open innovation: The new imperative for creating and profiting from technology*, Harvard Business Press.
- COCKBURN, I., R. HENDERSON, AND S. STERN (2019): “The Impact of Artificial Intelligence on Innovation,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.
- DAVID, P. A. (1990): “The dynamo and the computer: an historical perspective on the modern productivity paradox,” *The American Economic Review*, 80, 355–361.
- DIXIT, A. K., R. K. DIXIT, AND R. S. PINDYCK (1994): *Investment under uncertainty*, Princeton university press.
- FLEMING, L. (2001): “Recombinant uncertainty in technological search,” *Management science*, 47, 117–132.
- FLEMING, L. AND O. SORENSON (2004): “Science as a map in technological search,” *Strategic Management Journal*, 25, 909–928.
- GAVETTI, G. AND D. LEVINTHAL (2000): “Looking forward and looking backward: Cognitive and experiential search,” *Administrative science quarterly*, 45, 113–137.
- GOH, G. B., N. O. HODAS, AND A. VISHNU (2017): “Deep learning for computational chemistry,” *Journal of computational chemistry*, 38, 1291–1307.
- GOMES, J., B. RAMSUNDAR, E. N. FEINBERG, AND V. S. PANDE (2017): “Atomic convolutional networks for predicting protein-ligand binding affinity,” *arXiv preprint arXiv:1703.10603*.
- GROSSMAN, G. AND E. HELPMAN (1991): “Quality Ladders and Product Cycles,” *Quarterly Journal of Economics*, 106, 557–586.
- JONES, C. (1995): “R&D-Based Models of Economic Growth,” *Journal of Political Economy*, 103, 759–784.

- JONES, C. I. (2005): “Growth and ideas,” in *Handbook of economic growth*, Elsevier, vol. 1, 1063–1111.
- KAUFFMAN, S. (1993): *The origins of order: Self-organization and selection in evolution*, Oxford University Press, Oxford and New York.
- KAUFFMAN, S., J. LOBO, AND W. G. MACREADY (2000): “Optimal search on a technology landscape,” *Journal of Economic Behavior & Organization*, 43, 141–166.
- KLEINBERG, J., J. LUDWIG, S. MULLAINATHAN, AND Z. OBERMEYER (2015): “Prediction policy problems,” *American Economic Review*, 105, 491–95.
- LECUN, Y., Y. BENGIO, AND G. HINTON (2015): “Deep learning,” *Nature*, 521, 436.
- LEUNG, M. K., A. DELONG, B. ALIPANAHI, AND B. J. FREY (2016): “Machine learning in genomic medicine: a review of computational problems and data sets,” *Proceedings of the IEEE*, 104, 176–197.
- LEVINTHAL, D. A. (1997): “Adaptation on rugged landscapes,” *Management science*, 43, 934–950.
- MCCALL, J. J. (1965): “The economics of information and optimal stopping rules,” *The Journal of Business*, 38, 300–317.
- MULLAINATHAN, S. AND J. SPIESS (2017): “Machine learning: an applied econometric approach,” *Journal of Economic Perspectives*, 31, 87–106.
- NELSON, R. R. AND S. G. WINTER (1982): *An evolutionary theory of economic change*, Cambridge, MA: Harvard University Press.
- PISANO, G. P. (2006): *Science business: The promise, the reality, and the future of biotech*, Harvard Business Press.
- POPPER, K. R. (1959): *The logic of scientific discovery*, Hutchinson: London.
- PYZER-KNAPP, E. O., K. LI, AND A. ASPURU-GUZZIK (2015): “Learning from the harvard clean energy project: The use of neural networks to accelerate materials discovery,” *Advanced Functional Materials*, 25, 6495–6502.
- RACCUGLIA, P., K. C. ELBERT, P. D. F. ADLER, C. FALK, M. B. WENNY, A. MOLLO, M. ZELLER, S. A. FRIEDLER, J. SCHRIER, AND A. J. NORQUIST (2016): “Machine-learning-assisted materials discovery using failed experiments,” *Nature*, 533, 73 EP –.
- RIVKIN, J. W. (2000): “Imitation of complex strategies,” *Management science*, 46, 824–844.
- ROBERTS, K. AND M. L. WEITZMAN (1981): “Funding criteria for research, development, and exploration projects,” *Econometrica: Journal of the Econometric Society*, 1261–1288.
- ROMER, P. (1990): “Endogenous Technical Change,” *Journal of Political Economy*, 94, S71–S102.
- (1993): “Two Strategies for Economic Development: Using and Producing Ideas,” *Proceedings of the World Bank Annual Conference on Development Economics*.

- SCHUMPETER, J. A. (1939): *Business cycles*, vol. 1, McGraw-Hill New York.
- SHWARTZ-ZIV, R. AND N. TISHBY (2017): “Opening the black box of deep neural networks via information,” *arXiv preprint arXiv:1703.00810*.
- SPARKES, A., W. AUBREY, E. BYRNE, A. CLARE, M. N. KHAN, M. LIAKATA, M. MARKHAM, J. ROWLAND, L. N. SOLDATOVA, K. E. WHELAN, M. YOUNG, AND R. D. KING (2010): “Towards Robot Scientists for autonomous scientific discovery,” *Automated Experimentation*, 2, 1.
- STIGLER, G. J. (1961): “The economics of information,” *Journal of political economy*, 69, 213–225.
- TABOR, D. P., L. M. ROCH, S. K. SAIKIN, C. KREISBECK, D. SHEBERLA, J. H. MONTROYA, S. DWARAKNATH, M. AYKOL, C. ORTIZ, H. TRIBUKAIT, ET AL. (2018): “Accelerating the discovery of materials for clean energy in the era of smart automation,” *Nat. Rev. Mater.*, 3, 5–20.
- TADDY, M. (2019): “The Technological Elements of Artificial Intelligence,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.
- TEODORIDIS, F. (2018): “Understanding Team Knowledge Production: The Interrelated Roles of Technology and Expertise,” *Management Science*, 64, 3625–3648.
- TJUR, T. (2009): “Coefficients of determination in logistic regression models—A new proposal: The coefficient of discrimination,” *The American Statistician*, 63, 366–372.
- USHER, A. P. (1929): *A history of mechanical inventions: revised edition*, New York: McGraw-Hill.
- VARIAN, H. (2019): “Artificial intelligence, economics, and industrial organization,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by Agrawal, Gans, and Goldfarb, University of Chicago Press: Chicago, IL.
- VIRSHUP, A. M., J. CONTRERAS-GARCÍA, P. WIPF, W. YANG, AND D. N. BERATAN (2013): “Stochastic voyages into uncharted chemical space produce a representative library of all possible drug-like compounds,” *Journal of the American Chemical Society*, 135, 7296–7303.
- WAINBERG, M., D. MERICO, A. DELONG, AND B. J. FREY (2018): “Deep learning in biomedicine,” *Nature biotechnology*, 36, 829.
- WALLACH, I., M. DZAMBA, AND A. HEIFETS (2015): “AtomNet: A deep convolutional neural network for bioactivity prediction in structure-based drug discovery,” *arXiv preprint arXiv:1510.02855*.
- WALLACH, I. AND A. HEIFETS (2018): “Most Ligand-Based Classification Benchmarks Reward Memorization Rather than Generalization,” *Journal of Chemical Information and Modeling*, 58, 916–932, PMID: 29698607.
- WEITZMAN, M. L. (1979): “Optimal search for the best alternative,” *Econometrica: Journal of the Econometric Society*, 641–654.

——— (1998): “Recombinant growth,” *The Quarterly Journal of Economics*, 113, 331–360.

WRIGHT, S. (1932): “The roles of mutation, inbreeding, crossbreeding and selection in evolution,”
in *Proceedings of the sixth international congress of Genetics*, vol. 1, 356–366.

Figure 1: Unit Step Function Representing the Ranking Function for the Ground Truth

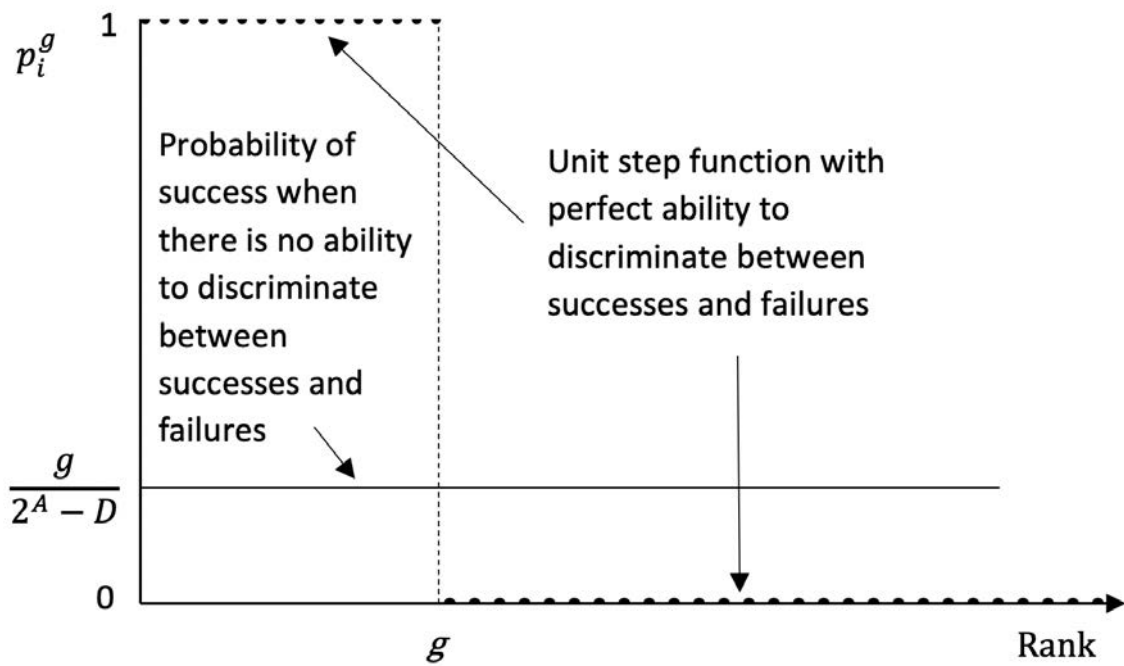


Figure 2: Ranking Function Curves for Different Values of the Discrimination Parameter, b






Figure 3:




(a) Determination of the Optimal Number of Tests




(b) Impact of an Improvement in the Prediction Model on the Optimal Number of Tests


Figure 4:



(a) Determination of the Optimal Number of Tests when the Innovator has a Single Innovation Target



(b) Impact of an Improvement in the Prediction Model when the Innovator has a Single Innovation Target and the Crossover Probability is below c



(c) Impact of an Improvement in the Prediction Model when the Innovator has a Single Innovation Target Crossover Probability is above c

Figure 5: Generic Workflow of the Two-Task Model

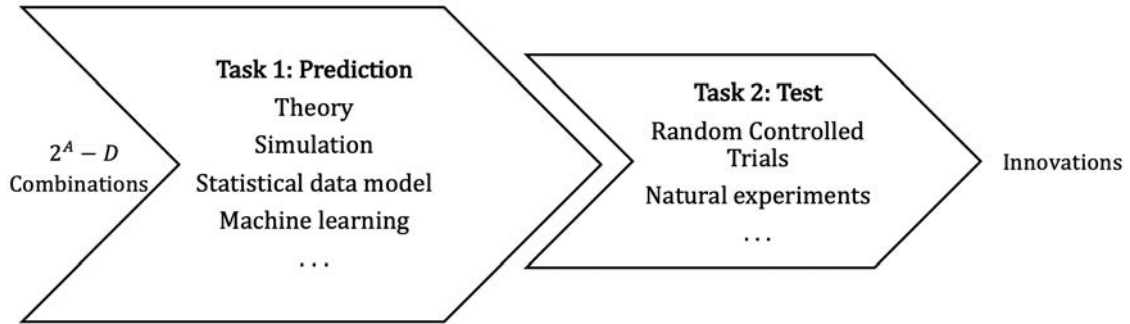


Figure 6: A Five-Stage Discovery Process with an Initial Prediction Stage and a Final Test Stage and Three Intermediate Prediction Refinement/Preliminary Testing Stages

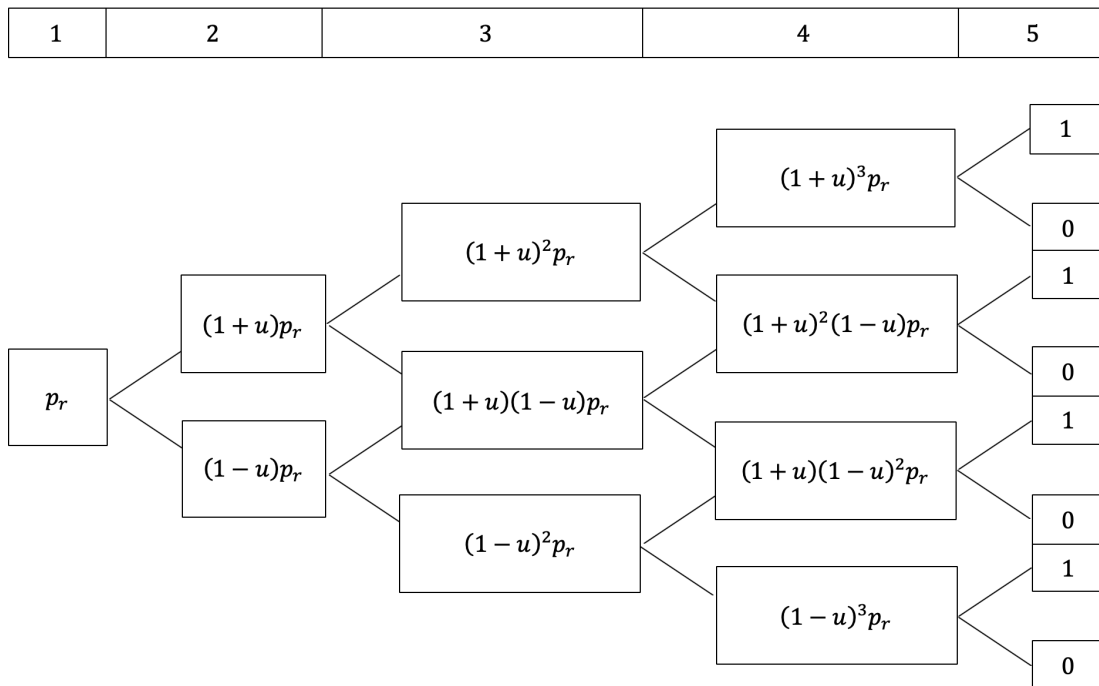


Figure 7: Expected Marginal Gross Value and Marginal Cost Curves with and without an Option to Abandon for the Three-Stage Discovery Process

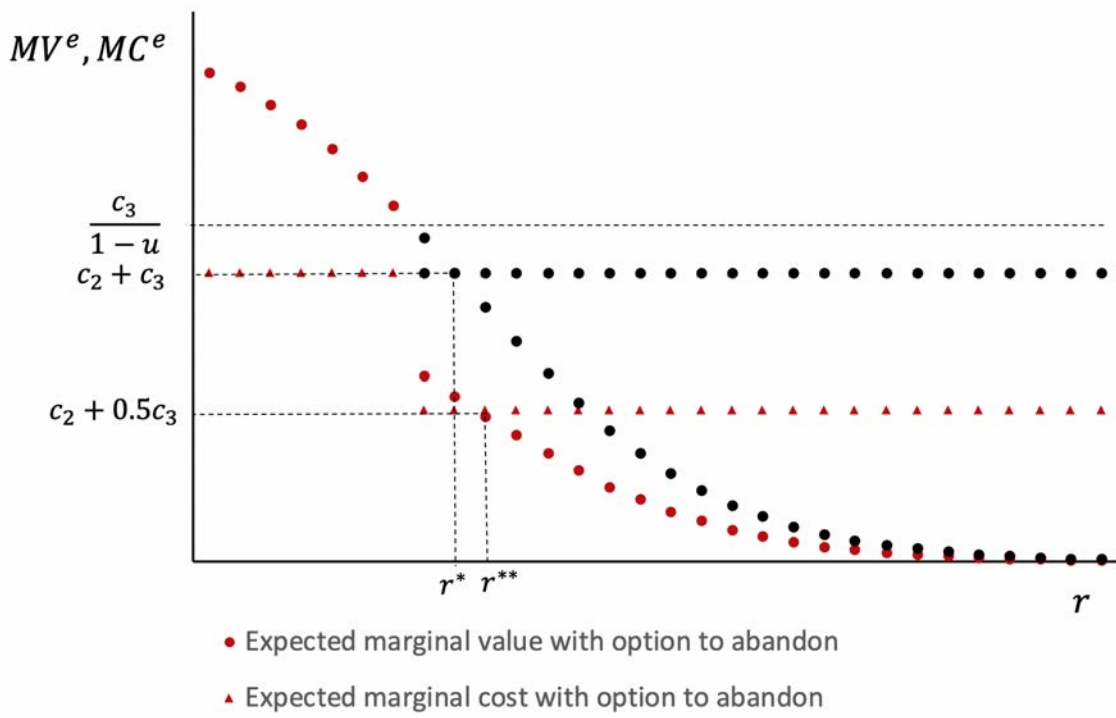
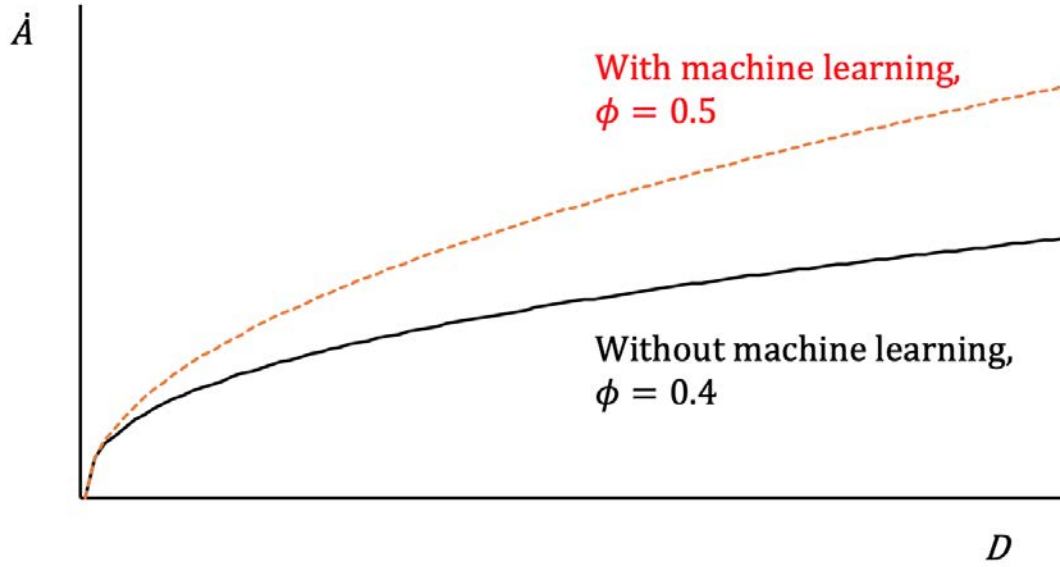



Figure 8:



(a) Relationship between Data and Innovation Performance



(b) Relationship between Idea Stock and Innovation Performance