International Journal of Industrial Organization 27 (2009) 449–462

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Disclosure or secrecy? The dynamics of Open Science[☆]

Arijit Mukherjee [a], Scott Stern [b,*]

[a] Bates White, LLC and Michigan State University, United States
[b] Kellogg Graduate School of Management, Northwestern University, United States

ARTICLE INFO

Article history:
Received 8 November 2007
Received in revised form 19 November 2008
Accepted 30 November 2008
Available online 6 December 2008

JEL classification:
O31
O32
L33

Keywords:
Open Science
Pasteur's Quadrant
Overlapping generations

ABSTRACT

Open Science is a dynamic system of knowledge production that depends on the disclosure of knowledge by researchers as an input into knowledge production by future researchers. To analyze the conditions supporting Open Science, we develop an overlapping generations model that focuses on the trade-off between disclosure and secrecy. While secrecy yields private returns that are independent of the actions of future generations, the benefits of disclosure depend in part on the use of disclosed knowledge by the subsequent researchers. We show that Open Science and Secrecy are both potential equilibria, and that the feasibility of Open Science depends on factors such as the costs of accessing knowledge from prior generations and the relative benefits to private exploitation under secrecy versus disclosure. In parameter regions where both Open Science and Secrecy can be supported, Open Science is associated with a higher level of social welfare. The analysis has policy implications for a number of areas, including public support for research training, appropriate design of formal intellectual property, and the role of scientific norms and institutions (such as an effective peer review process) in maintaining Open Science over the long run.

1. Introduction

At least since the development of scientific societies and related research institutions in the seventeenth century, the centrality of cumulativeness in scientific and technical advance has been recognized, most famously by Newton, who observed that scientific progress depends on “standing on the shoulders of giants.” While economic theory has focused on deriving the implications of cumulativeness for related economic variables such as the equilibrium growth rate (Romer, 1990; Grossman and Helpman, 1991; Jones, 1995; Jones, forthcoming) or the incentives for commercial innovation (Scotchmer, 1991; Gallini and Scotchmer, 2002; Scotchmer, 2004), relatively little research has focused on the microeconomic conditions that support a cumulative research environment.

The fact that knowledge is produced does not guarantee that follow-on researchers will be able to exploit that knowledge (Polanyi, 1967). Effective diffusion of knowledge across researchers and over time requires that individuals are aware of the extant knowledge and that they pay the costs of accessing that knowledge. The ability of a society to stand on the shoulders of giants depends not only the amount of knowledge it generates but also on the quality of mechanisms for storing knowledge, the trustworthiness of that knowledge, and the cost to future generations of accessing that knowledge (Mokyr, 2002; Furman and Stern, 2008).

Open Science is perhaps the most well-known system for achieving these objectives. Open Science is characterized by a distinctive set of incentives for cumulative knowledge production, including norms that facilitate disclosure and knowledge diffusion (Merton, 1973; Dasgupta and David, 1994). This system includes the recognition of scientific priority by future scientific generations, the importance of demonstrating experimental replicability, and a system of public (or coordinated) expenditures to reward those who contribute to cumulative knowledge production over the long term. By conditioning career rewards (such as tenure) on disclosure through publication, Open Science promotes cumulative discovery.[1] However, the logic underlying Open Science as an economic institution is more subtle. The ability to sustain disclosure over time depends not simply on the willingness of scientists to invest in research per se but also in their willingness to (1) invest in drawing upon the knowledge

☆ We thank Joshua Gans, Maria Goltsman, Tapas Kundu, Fiona Murray, Marcin Peski, Luis Vasconcelos, Michael Whinston, the Editor, two anonymous referees, and especially Julien Jamison for extremely helpful comments and suggestions. The first author gratefully acknowledges financial support from Bates White Research Enhancement Grant. The second author gratefully acknowledges funding from NSF Science of Science Policy Grant 0738394. All errors that may remain are ours.

* Corresponding author.
E-mail addresses: arijit@msu.edu (A. Mukherjee), s-stern2@northwestern.edu (S. Stern).

[1] Indeed, the policy rationale for public support of Open Science has been rooted in the idea that basic research is a public good and that ensuring an appropriate level of basic research requires some form of subsidy, most likely provided by the public sector (Nelson, 1959; Arrow, 1962). David (2004) builds upon and then extends a rich literature in the history of Science to emphasize that Open Science has long relied on the politically motivated patronage of key individuals. It is only within the past century that national governments have taken the lead in providing stable and extensive funding for Open Science (Nelson and Rosenberg, 1994; Mowery and Rosenberg, 1998).

0167-7187/$ – see front matter
doi:10.1016/j.ijindorg.2008.11.005

provided by prior researchers and (2) disclose their own discoveries in a way that can be accessed and exploited by future researchers.[2]

The ability to maintain Open Science may be challenged when discoveries are not only of scientific interest but also have significant commercial application. When a single discovery has *dual* applications – it can serve as an input to future scientific research and be exploited directly for commercial gain – a trade-off arises between the incentives to disclose through the scientific literature and the incentives to maximize direct commercial exploitation (Rosenberg, 1990; Stokes, 1997). Consider the Oncomouse (Murray, 2006). In the early 1980s, Professor Phil Leder at the Harvard Medical School developed the first genetically engineered mouse; it was called the Oncomouse. Leder and his colleague had used newly emerging transgenic techniques to insert an oncogene into a mouse embryo; the result was a mouse that was highly susceptible to cancer. Using the mice to examine the importance of oncogenes in the onset of cancer, Leder came to recognize that "it could serve a variety of different purposes, some purely scientific others highly practical" (Kevles, 2002, p. 83). This research was published in *Cell* in 1984, and, in 1988, a broad patent for the Oncomouse was granted by the USPTO. Harvard's licensee DuPont aggressively enforced these rights, including demands for "reach-through" rights and review of publications that used the Oncomouse in further scientific research. Over the next decade, a number of controversies surrounded the access to and credit for discoveries based on the Oncomouse. The conflict over the Oncomouse centered on the ability of the broader scientific community to exploit the Oncomouse (and to provide informal recognition to Leder and his coauthors) versus the incentives of DuPont to limit the diffusion of the Oncomouse in order to maximize its commercial advantage (Murray, 2006).

Although traditional models of science and innovation have often assumed a sharp delineation between purely scientific research and commercial applications, qualitative studies of scientific research have increasingly emphasized the importance of dual-use research (Rosenberg, 1974; Stokes, 1997; Murray, 2002; Murray and Stern, 2007). Stokes, in particular, suggested that a significant share of all scientific research combines the scientific and commercial motives and results in knowledge production in "Pasteur's Quadrant."[3] Pasteur's fundamental insights into microbiology simultaneously had practical applications for cholera and rabies while also serving as the foundation for the germ theory of disease (Geison 1995; Stokes, 1997).

This paper analyzes the feasibility of Open Science when research is conducted in Pasteur's Quadrant (i.e., has both scientific and commercial importance). We consider how incentives for access to prior knowledge, investment in knowledge, and the disclosure of discoveries depend on the disclosure and investment decisions of prior researchers and the access decisions of future researchers. A critical ingredient of our analysis is the fact that the incentives of any one researcher to participate in Open Science depend crucially on the choices of other researchers — i.e., the incentives to publish research in an academic journal depend on future researchers building on that discovery and providing appropriate citations to it in their own research. We model scientific disclosure as an endogenous economic outcome of the microeconomic environment, with the potential for Open Science depending on strategic interaction among researchers in their access, investment, and disclosure decisions.

Our model highlights two features of Open Science: (1) the ability to draw upon prior (disclosed) research and (2) the fact that the incentives to produce and disclose abstract knowledge depend on receiving credit from follow-on researchers. In contrast, the incentives for commercially motivated knowledge production are premised on the ability to limit the use of knowledge by others; we call this approach "Secrecy." Of course, the private returns to scientific research crucially depend on several exogenous factors such as the institutional and legal environment of the time. In our model, this is achieved through trade secrecy. (We consider the role of formal intellectual property rights (IPR) in an extension.) We embed the choice between secrecy versus disclosure into an overlapping generations framework in which each generation is composed of a single researcher who lives for two periods. During his first period of life, each researcher produces a knowledge output by choosing (1) whether to draw upon knowledge (if available) produced by the previous generation, (2) the level of investment in his own research, and (3) whether to disclose the produced knowledge for follow-on researchers in the next period. Each researcher faces a fixed cost of drawing upon prior knowledge, and a constant marginal cost of investment in his own research. The benefits to each researcher are composed of (1) the benefits from citations to his research by the next generation (if he chooses to disclose, and the next generation chooses to build on that research) and (2) private rents from proprietary exploitation of his knowledge. Researchers face a trade-off between maximizing the benefits from *private* exploitation (through secrecy) and earning a lower benefit from private exploitation but earning additional benefits from disclosure through the institutions of Open Science.

We draw out the equilibrium implications of this choice between secrecy and disclosure and focus on three potential outcomes: (1) "Open Science," in which each generation invests in access to prior knowledge, chooses a constant level of investment, and discloses knowledge to the next generation; (2) "Secrecy," in which each generation does *not* build on the knowledge produced by the prior generation, chooses a constant level of investment, and chooses not to disclose the knowledge produced to the subsequent generation; and (3) k-period "cycle" equilibrium, in which a single period of "Secrecy" is followed by $k-1$ periods of "Open Science."

At least one of these three types of equilibria *must* exist for any set of parameter values that describes the microeconomic environment. With that said, the feasibility of a given equilibrium depends crucially on the parameters of the economic environment. For example, the viability of Open Science is decreasing in both the cost of accessing knowledge produced by prior generations and in the relative benefits to private exploitation under secrecy versus disclosure. We also examine the role of factors such as the effectiveness of scientific institutions in promoting the effective transfer of knowledge across generations and the marginal cost of research investment. Rather than being grounded in differences in the type of knowledge produced, the model suggests that the feasibility of Open Science depends on the institutional and microeconomic environment in which that knowledge is produced; these parameters are themselves functions of the policy environment.

The model also highlights the potential for multiple equilibria for a given set of parameters, so that the choice between "Open Science" and "Secrecy" is endogenous to the strategic interaction among researchers. When multiple equilibria exist, we are able to rank welfare. Open Science, whenever viable, generates more surplus than any regime involving Secrecy. Moreover, among the set of Open Science equilibria, welfare increases as a function of the level of research investment. Finally, we considers a number of extensions and implications of the model: (1) the potential for knowledge spillovers across multiple generations (relaxing our assumption in the baseline model that spillovers only occur across immediately adjacent research generations), (2) the potential for hysteresis (is it more difficult to establish Open Science as an equilibrium than to maintain that equilibrium once it is established?), and (3) the role of formal intellectual property rights such as patents. The contribution of this paper is to isolate the equilibrium implications of the trade-off that arises for each research generation between secrecy and disclosure

[2] As has been emphasized by, among others, Blumenthal et al (1997), scientific researchers often withhold key materials or tools from follow-on researchers. This results in increasing policy concerns over access and transparency in the scientific commons.

[3] As in contrast with the knowledge produced for fundamental scientific interest (referred to as "Bohr's Quadrant") and the knowledge produced primarily for commercial gain (referred to as "Edison's Quadrant").

and to assess the implications of this equilibrium for welfare. By doing so, we contribute to two rapidly emerging literatures. First, building on Dasgupta and David (1994), several recent papers focus on the microeconomic conditions supporting "Open Science" as an economic institution (among others, Stern, 2004; Aghion et al., 2005; Lacetera, 2008; and Gans et al., 2008). At the same time, an emerging literature focuses on the incentives for knowledge disclosure by firms and on the interaction between trade secrecy and other mechanisms for earning returns from research investments (Horstmann et al., 1985; Arora, 1995; Anton and Yao, 2004; Lerner and Tirole, 2005; Kultti et al., 2007). This paper complements these contributions by focusing on the strategic impact of disclosure when the returns from knowledge production accrue both from citations from follow-on researchers and from traditional commercial returns.

The remainder of the paper is organized as follows. The next section introduces the basic model structure. Section 3 derives the equilibrium of the model and discusses how the set of equilibria may change in different economic environments. Section 4 considers several short extensions to the baseline model. A final section concludes. All proofs are in the Appendix A.

2. The model

We consider an overlapping generations framework. In each generation, a single researcher is born, and each generation lives for two periods. There is an infinite sequence of researchers; letting the tth generation be denoted as G_t $(t=0,\pm1,\pm2,\ldots)$, the generations G_{t-1} (currently in their second period of life) and G_t (currently in their first period of life) therefore coexist.

At the beginning of the first period of life, researcher G_t makes three choices: (1) whether to invest in accessing knowledge from G_{t-1}, (2) the level of current research investment, and (3) whether to allow the knowledge produced by G_t to be accessible to the next generation. The first of these decisions involves a sunk cost, F, which allows G_t to build directly on the publicly disclosed "knowledge pool" of the (currently living) prior research generation, G_{t-1}. To build upon prior research findings, researchers must often learn and master particular methods or research tools or engage in costly replication experiments. These efforts are investments that, unless they contravene prior findings, do not contribute directly to the knowledge pool (Mokyr, 2002; Furman and Stern, 2008). The sunk cost F is simply the cost for a researcher to reach the publicly disclosed knowledge frontier. While we do not separately model the researcher's investment in "general" knowledge or a baseline requirement for a researcher to produce knowledge at all (as in Jones, forthcoming), our analysis focuses on whether a researcher undertakes the investment that is required to be able to exploit frontier advances in a particular scientific research field. Let z_t be an index of the size/quality of the frontier knowledge potentially available to G_t, and let the indicator function a_t be equal to 1 if G_t accesses the frontier, and 0 otherwise.

Second, G_t chooses a level of research investment, x_t at a constant marginal cost ψ and zero fixed cost. The output of this investment is a new piece of knowledge produced at the end of the first period of his life, $\Delta K_t(x_t)$ that is (weakly) increasing in x_t. The private return to ΔK_t is equal to a sum of two streams of benefits. First, G_t can earn utility from citations to his research by the next generation (the number of citations equals C_t). Second, G_t can earn private rents from proprietary exploitation of that knowledge (the value of this private rent equals P_t). For example, G_t might seek to earn a commercial advantage from the knowledge through trade secrecy. (We briefly discuss the potential for formal intellectual property rights such as patents in Section 4.) Together, the payoff for G_t is

$$C_t + P_t - Fa_t - \psi x_t, \tag{1}$$

and it is realized at the end of the second period of G_t's life. In other words, similar to the duality emphasized by Stokes (1997), a single piece of knowledge can simultaneously be useful (potentially) for future research generations and in the context of private exploitation.

Third, G_t chooses whether to *disclose* his knowledge and make it potentially accessible to G_{t+1}. Let d_t be equal to 1 if G_t deposits the knowledge for public use, and 0 otherwise. The value from citations is equal to $\mu^C(a_t,x_t,d_t;z_t,a_{t+1})$, which represents the returns to G_t of access by G_{t+1}, for a given level of $\Delta K_t(x_t)$. Analogously, the value from private exploitation is equal to $\mu^P(a_t,x_t,d_t;z_t)$.[4] Our analysis depends crucially on our assumptions about the benefit streams from these two functions under different disclosure regimes:

Assumption 1. If G_t chooses $d_t=0$, $\mu^C=0$.

In other words, if no knowledge is disclosed, no citations are received. As a result, the benefits are simply equal to $\mu^P(a_t,x_t,0;z_t)$.

Assumption 2. If G_t chooses $d_t=1$,

(a) $\mu^C=0$ if $x_t=0$ or $a_{t+1}=0$;
(b) $\mu^C>0$, $\forall\ x_t>0$ when $a_{t+1}=1$.

In order to receive a positive level of citations, one must invest a positive amount in research and the *next* generation must choose to access from the public knowledge pool. As a result, the benefits from citations under disclosure depend on the strategy of G_{t+1}.

Assumption 3. $\mu^P(a_t,x_t,0;z_t)=\lambda\mu^P(a_t,x_t,1;z_t)$, $\forall a_t,x_t,z_t$, where $1<\lambda<\infty$.

The private benefits from knowledge are *lower* with public disclosure. For simplicity, we assume that this discount is a constant factor for all levels of knowledge output.

Assumption 4. The functions μ^C and μ^P are twice continuously differentiable in x_t and have the following characteristics:

(a) Boundedness: $\mu^C\in[0,\bar{\mu}^C]$ and $\mu^P\in[0,\bar{\mu}^P]$.
(b) Monotonicity: For $i\in\{C,P\}$, $\frac{\Delta\mu^i}{\Delta a_t}\geq 0$, $\frac{\partial\mu^i}{\partial x_t}>0$, and $\frac{\partial\mu^i}{\partial z_t}>0$ if $\alpha_t=1$, but independent of z_t otherwise.
(c) Strict Concavity: $\frac{\partial^2\mu^i}{\partial x_t^2}<0$; $i\in\{C,P\}$.
(d) Complementarity: $\frac{\Delta}{\Delta a_t}\left(\frac{\partial\mu^i}{\partial x_t}\right)\geq 0$, and $\frac{\partial^2\mu^i}{\partial z_t\partial x_t}\geq 0$ $(i\in\{C,P\})$ *only* if $a_t=1$, but equal to zero otherwise.
(e) Positive Investment: $\frac{\partial}{\partial x_t}(\mu^C+\mu^P)>\psi$ *at* $x_t=0$, $\forall(a_t, z_t, a_{t+1})$.

Both the citation and the private benefit functions are bounded above, and both functions are (weakly) increasing all of their arguments. Moreover, we assume that there are diminishing marginal returns to investment but that the returns to investment are (weakly) increasing in the quality of the knowledge pool (provided that G_t chooses to access the knowledge pool). Also, it is always worthwhile for G_t to invest a positive amount in research irrespective of the decisions of G_{t-1} and G_{t+1}. These assumptions ensure that there is a well-defined optimum, and they allow us to evaluate the impact of complementarity between access to the knowledge pool and the level of research investment.

The model highlights the relationship between research investment, disclosure choices and expectations about future decisions to access the knowledge pool. We assume that the quality/size of the knowledge pool to G_t, z_t depends on both the quantity of research deposited into the pool by the immediate prior research generation (i.e., ΔK_{t-1}) and the "quality" of the knowledge pool (parameterized by α). The ability to exploit knowledge across research generations depends on the quality of institutions that, by enhancing both the "technology of access" and the "trustworthiness of sources," facilitate low-cost knowledge transfer by enhancing both the "technology of access" and the "trustworthiness of sources" (Mokyr, 2002, p. 8). As

[4] The functions μ^C and μ^P can be interpreted as a reduced-form version of a stochastic citations function that is a random variable with a mean, $\tilde{\mu}^C(a_t,\Delta K_t(x_t),d_t;z_t,a_{t+1})$. Similarly, μ^P can be interpreted as a reduced-form version of $\tilde{\mu}^P(\alpha_t,\Delta K_t(x_t),d_t;z_t)$.

discussed further in Furman and Stern (2008), the quality of the knowledge pool depends on institutions that enhance the "fidelity" of the knowledge pool (i.e., the ability of researchers to trust and use prior research findings), such as the strength and consistency of the peer review process and the availability and clarity of data and resources for replication purposes.[5] We therefore interpret α as a measure of the institutional quality and fidelity of the knowledge pool that which enhances the efficiency by which knowledge from G_{t-1} can be exploited by G_t (conditional on deposit by G_{t-1}).

Implicit in the above formulation of the knowledge pool is the assumption that that useful learning for G_t only requires access to the knowledge generated by G_{t-1}. In other words, we are assuming that the *net* benefits of learning depend crucially on access to recently produced knowledge, either because the depreciation of knowledge is sufficiently high or because the relative benefits of access to "old" knowledge are relatively low. While this assumption is somewhat extreme in nature, it significantly improves the analytical tractability of the model. Moreover, some of the key insights of the model continues to hold under a more general setting where z_t depends on the knowledge produced by multiple prior generations.[6]

Together, these specifications can be captured by the following functional form of z_t.

$$z_t = d_{t-1}\tilde{z}(\Delta K_{t-1}(x_{t-1});\alpha) = d_{t-1}z(x_{t-1};\alpha). \tag{2}$$

We make the following assumption on the z function.

Assumption 5. The z function has the following properties: (a) z is a non-negative bounded function, and (b) z is increasing in all of its arguments.

Finally, although this assumption is not crucial, we assume zero discounting across periods within a generation's life.

3. Equilibrium secrecy or disclosure behavior

3.1. The game

As stated above, the economy is composed of overlapping generations G_t ($t=0,\pm1,\pm2,\ldots$). Each generation seeks to maximize its individual payoff:

$$U_t(a_t,x_t,d_t;z_t,a_{t+1}) = \mu^C(a_t,\ x_t,\ d_t;\ z_t,\ a_{t+1}) + \mu^P(a_t,\ x_t,\ d_t;\ z_t) - \psi x_t - a_t F \tag{3}$$

by choosing a triplet (a_t,x_t,d_t) conditional on the observed value of z_t and the strategy of G_{t+1}. Since d_t and a_t are binary, it is useful to (1) evaluate the optimal value of x_t, conditional on each potential combination of d_t and a_t, and (2) choose the triplet that yields the highest value for U_t. The form of the first-order conditions for x_t depends on the deposit decision, d_t. On the one hand, for $d_t=1$ and $a_t\in\{0,1\}$, G_t chooses x_t to solve:

$$\frac{\partial}{\partial x_t}\mu^C(a_t,x_t,1;\ z_t,\ a_{t+1}) + \frac{\partial}{\partial x_t}\mu^P(a_t,x_t,1;\ z_t) = \psi. \tag{4}$$

On the other hand, for $d_t=0$ and $\alpha_t\in\{0,1\}$, the optimal x_t solves:

$$\frac{\partial}{\partial x_t}\mu^P(a_t,x_t,0;\ z_t)\equiv\lambda\frac{\partial}{\partial x_t}\mu^P(a_t,x_t,1;\ z_t) = \psi. \tag{5}$$

Under disclosure, the marginal cost of research investment is set equal to the marginal benefits associated with enhancements to the citation stream as well as the marginal benefit of private exploitation. Under nondisclosure, the marginal cost of research is equated to the value arising from private exploitation. In other words, the marginal value of research investment within Pasteur's Quadrant depends crucially on (1) whether each research generation builds on prior knowledge, and (2) whether each research generation is a source of knowledge for subsequent generations. By considering each of the four combinations implied by Eqs. (4) and (5), G_t maximizes U_t for any realization of the choices of G_{t-1} and the strategy of G_{t+1}.

We denote $x_{a_t}^{d_t}$ as the optimal value of x_t given the values of a_t and d_t. Of course, the optimal value of x_t will depend on z_t and a_{t+1}, but we suppress them for the sake of brevity (unless the values of z_t and a_{t+1} are of any particular significance).

For the sake of analytical tractability, we restrict attention to a special class of Subgame Perfect Nash Equilibria (SPNE), namely, the pure strategy *periodic* SPNE. A pure strategy (for generation G_t) is a function $\sigma_t:\mathbb{R}_+\rightarrow\mathbb{R}_+\times\{0,1\}^2$ that maps the level of z_t into a triplet of decision variables (a_t,x_t,d_t), and so an SPNE is a sequence $\{\sigma_t\}_{t=-\infty}^{\infty}$ such that $\forall t$, σ_t maximizes U_t, given any z_t and σ_{t+1}. In other words, for any history of actions up to period $t-1$, G_t maximizes U_t, conditional on the (subgame perfect) strategy of G_{t+1}. A pure strategy SPNE is said to be a k-period *periodic* SPNE if on the equilibrium path $\{a_t,x_{a_t}^{d_t},d_t\}=\{a_{t+k},x_{a_{t+k}}^{d_{t+k}},d_{t+k}\}\forall t$.

While we cannot rule out the possibility that there might be other equilibria of this game that are not periodic, the class of periodic equilibria is, however, rich enough to encompass two types of equilibria that we are particularly interested in: the stable path equilibria and the cyclic equilibria. The stable path equilibria is equivalent to the case where $k=1$ (i.e., $\{a_t,x_{a_t}^{d_t},d_t\}_{t=-\infty}^{\infty}=\{a^*,x^*,d^*\}\forall t$). There are two broad types of stable path equilibria in this game: secrecy (i.e., nondisclosure) and disclosure. Under a stable nondisclosure regime, behavior along the equilibrium path is equal to $\{0,x_0^0,0\}\forall t$; under a stable disclosure regime, equilibrium behavior is equal to $\{1,x_1^1,1\}\forall t$. The cyclic equilibria are a purely periodic SPNE that involve a k-period cycle in which a single period of nondisclosure is followed by $k-1$ periods of disclosure.

Finally, it is worth noting that one can also potentially consider the mixed strategy equilibria of this game where G_t randomizes over his access (a_t) and disclosure decisions (d_t) and chooses the level of research investment, x_t based on the realized value of a_t and d_t. In such a scenario, given the value of z_t, G_t chooses a probability of access, say $p_t^a(z_t)$, a probability of disclosure, say $p_t^d(z_t)$, and an investment function $x_t^*(a_t,d_t;z_t)$ such that $\{p_t^a,p_t^d,x_t^*\}$ maximizes G_t's expected profit given the strategy of G_{t-1}. Analysis of mixed strategy equilibria would be considerably less tractable. Abstracting away from the (somewhat subtle) conditions that guarantee existence of a generic mixed strategy equilibria an infinitely repeated overlapping generations model (see, for example, Gossner, 1996), a *stable path* mixed strategy equilibrium would require the existence of a "fixed point" in the class of functions $\{p_t^a,p_t^d,x_t^*\}$. It is unclear whether this can be guaranteed under the current structure of the model.

3.2. The equilibrium

Our first task is to characterize the nature of the set of equilibria for this game. We begin by demonstrating that, for all parameter values, a periodic equilibrium exists. We then focus on those regions of the parameter space where stable path equilibria exist, including those regions where both disclosure and nondisclosure equilibria exist simultaneously. Our results highlight the nature of strategic interaction across research generations and the economic conditions

[5] The effectiveness of knowledge transfer across generations depends both on the transparency of disclosure by individual researchers (i.e., each generation) and on the institutional environment for cumulative knowledge production (Mokyr, 2002). To focus on the impact of policy and institutions in enhancing the efficiency of exploitation (as we discussed from an empirical perspective in Stern (2004) and Furman and Stern (2008), we model α as a parameter that is not directly influenced by the endogenous choices of each research generation but instead represents (external) policy or institutional changes.

[6] We discuss the consequences of relaxing this assumption in Section 4.

required to support "Open Science" as an equilibrium in an environment similar to that emphasized by Stokes (1997). Our first proposition demonstrates that, for any parameter value for this game, a periodic equilibrium exists.

Proposition 1. A periodic pure strategy SPNE exists for all parameter values. Moreover, each of the two stable-path equilibria (disclosure and nondisclosure) exists within certain parameter regions.

The existence of a periodic pure strategy SPNE follows from an analysis of the potential for stable path SPNE and from a demonstration that that a multiperiod pure strategy SPNE must exist if a stable path SPNE does not exist.[7] Specifically, we show that a nondisclosure equilibrium exists for sufficiently high levels of F or λ and that a disclosure equilibrium exists as long as *both* F and λ are sufficiently small. In other words, as long as the private exploitation benefits from nondisclosure, or the costs of accessing knowledge are sufficiently high, no disclosure can be induced. However, when access costs are relatively low, and the benefits from purely private exploitation are modest or low, a pure strategy SPNE can be sustained.

The remainder of the proof demonstrates that, in those "intermediate" cases (where both disclosure equilibrium and nondisclosure equilibrium do *not* exist), a multiperiod periodic SPNE must exist. The intuition for the existence of a multi-period periodic SPNE is subtle but instructive. A multiperiod periodic SPNE may arise because we have not yet placed structure on the relative sensitivity of μ^C and μ^P to z_t. In particular, if μ^C is sufficiently sensitive to z_t (relative to the sensitivity of μ^P to z_t), then either a disclosure or nondisclosure equilibrium must exist. However, if this condition is not satisfied, it may be the case that while disclosure may be an optimal response for G_t for low levels of z_t (when disclosure is "rewarded" by access by G_{t+1}), there exists a sufficiently high level of z_t such that private exploitation becomes optimal. A multiperiod periodic equilibrium is characterized by a single period of nondisclosure in G_t, yielding $z_{t+1}=0$ for G_{t+1}. Under this low level of z, research incentives are modest (because of complementarity between z and x) but G_{t+1} has an optimal strategy to disclose to G_{t+2}. The incentives for research for G_{t+2} are increased (relative to G_{t+1}), yielding a still higher level of z for G_{t+3}. This process continues until z_{t+k}, for some large enough k, reaches a threshold, at which point private exploitation becomes optimal. In other words, a k-period periodic equilibrium is characterized by a single period of nondisclosure, followed by $k-1$ periods of disclosure and increasing levels of research investment by the "scientific community," followed by private appropriation in the kth period.

One of the main insights of the model is that the key parameters determine the feasibility of different equilibria. We illustrate these issues through a series of corollaries that highlight the "bounds" on the parameters that sustain the nondisclosure and disclosure equilibria.

Corollary 1. There exists a value of λ, say λ^{nd}, and a value of F depending on λ, say $F^{nd}(\lambda)$, such that a nondisclosure equilibria must exist for all (F,λ) if either $F \geq F^{nd}(\lambda)$ or $\lambda > \lambda^{nd}$.

In other words, if *either* the fixed cost of accessing prior knowledge or the benefit to private exploitation associated with nondisclosure is sufficiently high, then a nondisclosure equilibrium must exist. While the existence of a nondisclosure equilibrium under unfavorable conditions is not particularly surprising, Corollary 1 also highlights that a nondisclosure equilibrium need not exist if *both* the fixed cost of accessing knowledge (F) and the benefits from purely private exploitation (λ) are sufficiently small. Indeed, when *both* of these conditions hold, at least one disclosure equilibrium always exists.

[7] Importantly the existence of a stable path SPNE is not generic. Proposition 1 depends on all of our assumptions, including the complementarity between z_t and x_t.

Corollary 2. There exists a value of F, say F^d, and for all $F \leq F^d$, a value of λ depending on F, say $\lambda^d(F)$, such that disclosure equilibrium must exist for all $(F,\lambda) \leq (F^d,\lambda^d(F))$.

The existence of a disclosure equilibrium requires that both F and λ are sufficiently small.[8] In other words, even if there is only a very small benefit to purely private exploitation, but the costs of accessing knowledge across generations are prohibitively high, then each research generation will choose nondisclosure in the expectation that, had it disclosed, its knowledge would not have been accessed by the follow-on generation. Similarly, even if the fixed cost of accessing knowledge is extremely small, the feasibility of a disclosure equilibrium requires that the benefits from purely private exploitation should not be too large.

To fully characterize the feasibility of equilibria across different regions of interest, we have to further establish (1) those regions where *both* disclosure and nondisclosure equilibria may be feasible and (2) those regions where *neither* disclosure nor nondisclosure equilibria may be feasible. To do so, we first characterize the $\lambda^d(F)$ function. Consider the set of all disclosure equilibria. Each of these equilibria is associated with a constant value of z_t, and let Z^F be the set of all such z_t values.

Lemma 1. Given a value of $z \in Z^F$, for all $F < F^d$, there exists a value of λ, say $\lambda(z)$, and a function $\lambda^*(F,z)$ that is decreasing in F, such that $\lambda^d(F) = \sup_{z \in Z}\{\min\{\bar{\lambda}(z),\lambda^*(F,z)\}\}$. Moreover, for any $z \in Z^F$, $\lambda^*(F,z) \geq \lambda^{nd}$ for all $F < F^d$.

Lemma 1 raises the possibility that there are regions of overlap between the two types of stable path equilibria and regions in which no stable path equilibrium need exist. We now characterize each of these two environments. (We omit the formal proof of this result because it follows from Corollary 1 and 2 and the fact $\lambda^d(F)$ is decreasing in F (as $\lambda^*(F,z)$ is decreasing in F).

Corollary 3. If for any $F = F' \leq F^d$, $\lambda^d(F') \geq \lambda^{nd}$, then for all $F \leq F'$ and for all $\lambda \in [\lambda^{nd},\lambda^d(F)]$, both disclosure and nondisclosure equilibria coexist.

As long as F as is sufficiently low and λ is of an intermediate value, both types of stable path equilibria may exist. The intuition is as follows. G_t's access and disclosure decision are endogenous to the investment and disclosure choices of the prior generation and equilibrium access strategy of G_{t+1}. In particular, the incentives to access the knowledge of a prior generation is strictly increasing in z_t, which is the quality of knowledge pool accessible to G_t. As such, the access choice of G_t can be described by a decision rule to sink F when $z_t > z'$, where z' is determined by parameters such as F and λ, as well as the strategy of G_{t+1}. In the region described in Corollary 3, lack of disclosure in G_{t-1} reduces the productivity of research and investment incentives for G_t. As such, given the cost of research ψ, the optimal level of research (and the resulting pool quality z_t) are insufficiently large to induce access by G_{t+1}. Conversely, if G_{t-1} had disclosed, this would raise research productivity and incentives and result in a level of research output that is sufficient to induce G_{t+1} to sink F to access that knowledge if it were disclosed. In other words, when access costs are sufficiently low and the disclosure "penalty" (in terms of private exploitation) is at an intermediate range, *either* a secrecy regime *or* a disclosure regime can be supported, depending on the disclosure and investment behavior of prior generations and (equilibrium) access behavior of future generations.

[8] Recall that by assumption, $\lambda > 1$. Obviously, if $\lambda \leq 1$, disclosure can always be part of an "optimal" strategy, since there are no direct costs and no negative impact in terms of private exploitation. Indeed, when $\lambda < 1$, the value from private exploitation is increasing in disclosure—a condition, which may hold when, for example, a firm is attempting to popularize a novel technological standard, which also merits scientific interest.

It is important to note that among the potential equilibria in these regions, the level of research investment under disclosure need not be higher than the level of research investment under secrecy. Specifically, though the productivity of research for a given level of investment is higher under disclosure, the incentives for research depend on λ. If λ is sufficiently high (but not so high as to preclude a disclosure equilibrium), it is possible that the level of equilibrium research investment in the nondisclosure equilibrium is higher than the minimum level of investment that can sustain a disclosure equilibrium.

Finally, when $F^d < F^{nd}$ and/or $\lambda^d \leq \lambda^{nd}$ no stable path equilibrium may exist for some regions. In particular, if $F^d < F^{nd}$, then with $F \in (F^d, F^{nd})$ and λ sufficiently low, it is possible that the only feasible pure strategy periodic equilibrium has $k > 1$. Similarly, if $\lambda^d \leq \lambda^{nd}$, then for any $F < F^{nd}$ and $\lambda \in (\lambda^d, \lambda^{nd})$, no pure strategy stable path equilibrium may exist. As mentioned earlier, equilibrium behavior in this region consists of $k-1$ periods of disclosure and increasing levels of research investment, followed by a single period of secrecy. In other words, this model incorporates the possibility that, along the equilibrium path, private knowledge exploitation occurs only when the knowledge pool is sufficiently "rich."

Our discussion is summarized in Fig. 1, where we characterize different equilibria that exist for different values of F and λ. From a comparative statics perspective, the feasibility of disclosure is decreasing in both F and λ. In Region I, when either access costs or the penalty from disclosure (or both) is sufficiently high, the only pure strategy SPNE is composed exclusively of secrecy. On the other hand, in Region III, when *both* access costs and the penalty from disclosure are sufficiently low, disclosure is supported as a pure strategy equilibrium (and no pure secrecy equilibrium may be feasible). Region II describes those environments where both disclosure and secrecy may be feasible; Region IV describes the parameter values where the only pure strategy SPNE may be periodic.

3.3. Comparative statics

We now examine the impact of changes in the fidelity of disclosed scientific knowledge (α) and the cost of research (ψ) on equilibrium research investment, disclosure, and access choices. Our comparative statics analysis first evaluates the impact of these parameters on the level of research investment, conditional on a given access and disclosure "regime." Specifically, we investigate how the optimal research investment under secrecy (x_0^0) and disclosure (x_1^1) varies with α and ψ, where the change is small enough such that secrecy (or disclosure) continues to be an equilibrium of the game.

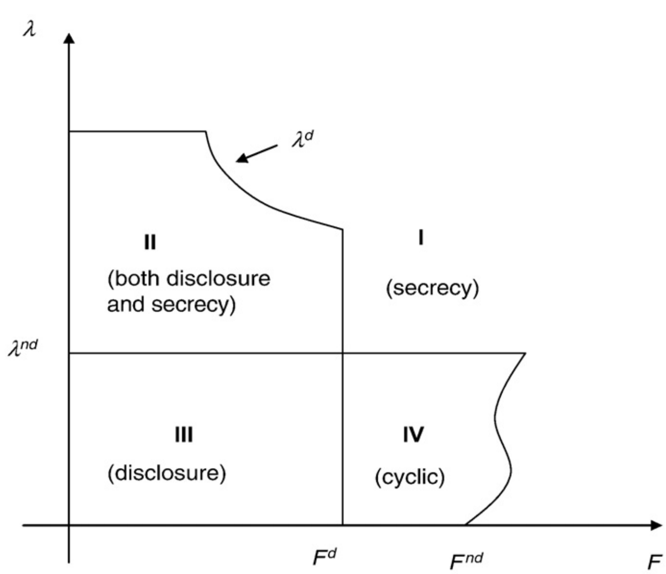

Fig. 1. Types of equilibria for different values of F and λ.

Proposition 2. (i) Both x_0^0 and x_1^1 are decreasing in ψ. (ii) x_0^0 is independent of α but x_1^1 is increasing in α.

On the one hand, an increase in ψ or a decrease in α results in a direct effect that stems from the concavity of research productivity in x: x must decrease because now the marginal benefit of research is lower than its marginal cost. As well, when considering the comparative static with respect to x_1^1, we must also consider an indirect effect, since a decrease in x_1^1 also reduces z, the size/quality of the knowledge pool. This indirect effect of a reduction in z also reduces the marginal benefit of research, and so it further reduces x_1^1 and reinforces the direct effect. Consequently, within both the secrecy and disclosure regimes, research investments are strictly decreasing in the cost of doing research but weakly increasing in the fidelity of the prior generation's knowledge. More generally, Proposition 2 highlights an important role for scientific research institutions: rather than simply enhancing research productivity, such institutions can raise equilibrium research incentives.

We next examine how the feasibility of disclosure and nondisclosure equilibria are affected by α and ψ. Earlier, we derived sufficient conditions for the existence of each of these two forms of equilibrium in terms of the boundary points F^{nd}, λ^{nd}, F^d and λ^d. By studying how these boundary points change with α and ψ, we gain insight how these parameters impact the equilibrium choice between disclosure and secrecy. In particular, the analysis demonstrates a sharp relationship between these boundary points and the fidelity of knowledge but a more ambiguous relationship between these boundary points and the costs of research investment.

Proposition 3a. (i) F^{nd} is increasing in α, (ii) λ^{nd} is independent of α, (iii) F^d is increasing in α, and (iv) λ^d is increasing in α if $\partial\mu^P/\partial z$ is sufficiently small compared to $\partial\mu^C/\partial z$, but ambiguous otherwise.

Recall that F^{nd} and λ^{nd} are the smallest values of F and λ *above* which a secrecy equilibrium is guaranteed to exist, and that F^d and λ^d are the largest values of F and λ *below* which a disclosure equilibrium is guaranteed to exist. As the fidelity of the knowledge stock increases, it becomes harder to sustain a secrecy equilibrium and may become easier to sustain a disclosure equilibrium (and this latter result is guaranteed when a change in the level of knowledge pool (z) has a more pronounced effect on the returns from citation (μ^C) than on the returns from private expropriation (μ^P)).

Consider F^{nd}, the minimum access cost at which accessing the knowledge pool becomes unprofitable for G_{t+1} even when G_t deviates from a secrecy equilibrium and deposits its knowledge. As α increases, the payoff of G_{t+1} from accessing the knowledge (deposited by G_t) increases. Thus, F^{nd} must increase to ensure that accessing the knowledge remains unprofitable.

Analogously, F^d is the cost of accessing the knowledge pool at which G_t is indifferent between the payoffs associated with a disclosure equilibrium and the payoff from the deviation where G_t does not access the knowledge pool, but continues to deposit his knowledge. Clearly, the equilibrium payoff under disclosure increases with α, as α increases z_t. However, the payoff from deviation remains unchanged since G_t does not access the knowledge pool under the deviation considered above. As a result, F^d must increase to maintain the indifference condition that defines this boundary.

The threshold λ^{nd} is the value of λ at which G_t is indifferent between the equilibrium payoff under secrecy and the payoff from disclosing knowledge when G_{t+1} accesses the knowledge pool. Because the knowledge pool is necessarily empty in a secrecy equilibrium, enhancing the fidelity of the knowledge pool impacts neither G_t's equilibrium payoff nor G_t's payoff from deviation (i.e., disclosure). Consequently, λ^{nd} is independent of α.

Finally, the analysis of the impact of α on λ^d is more subtle. λ^d is the largest value of λ at which G_t's payoff under disclosure is greater than the payoff from a deviation to nondisclosure (but under the assumption that G_t can also choose whether to continue to access the prior generation's knowledge or not). An increase in α increases *both* the payoff from disclosure and the payoff from a deviation towards nondisclosure. Since we cannot argue *a priori* which payoff increases more, the impact of a change in α on λ^d is ambiguous. However, if we further assume that the benefits from a high-fidelity knowledge pool are particularly important in the case where the benefits arise from citations by future researchers (i.e., a given change in z has a more pronounced effect on μ^C than on μ^P), then the payoff from continuing to disclosure increases more than the payoff from the deviation to secrecy. Consequently, under this further assumption, λ^d is also increasing with α.

Proposition 3b. The impact of ψ on the boundary points F^{nd},λ^d, and λ^{nd} is ambiguous, but F^d is decreasing in ψ.

The ambiguity of the relationship between the boundary points and the marginal cost of research is a consequence of the fact that the indifference condition determining each of the boundary points (expect F^d) involves a comparison of G_t's payoff across different disclosure decisions. For example, λ^{nd} is determined by a comparison of the payoff under secrecy and the payoff from a deviation to disclosure. Moreover, these comparisons depend themselves on the optimal level of research investment associated with a particular "regime," x_a^d; in other words, the boundary conditions are themselves sensitive to the level of research investment. The ambiguity stems from the fact that we cannot rank order the optimal levels of research investments across different disclosure regimes. However, in the case of F^d, we are unambiguously able to determine the relevant magnitudes, since this comparison depends on the payoff to G_t from disclosure for different access decisions. Using the fact that $x_1^1 > x_0^1$ (by complementarity between x_t and a_t), we can argue that F^d is decreasing in ψ.

3.4. Welfare analysis

As we noted earlier, there are significant parameter regions in which both disclosure and nondisclosure equilibria are feasible. In these regions, there may be multiple disclosure equilibria along with a unique nondisclosure equilibrium. While we cannot rule out multiple equilibria, we are able to offer a welfare comparison between these equilibria.

Proposition 4. If 'disclosure' (i.e. $d_t = 1$) and 'nondisclosure' (i.e. $d_t = 0$) stable path equilibria coexist, the researcher's payoff in the 'disclosure' equilibrium is strictly greater than his payoff in the 'nondisclosure' equilibrium.

In other words, welfare along the *minimal* disclosure equilibrium is (weakly) superior to welfare along the nondisclosure equilibrium, even when both are potential SPNE. It is useful to recall that the minimal level of research investment that sustains a disclosure equilibrium need not be higher than the unique level of research investment associated with a secrecy equilibrium. However, for a disclosure equilibrium to be sustainable, each research generation must prefer to remain in the disclosure equilibrium relative to "autarky," where a generation neither draws on prior knowledge nor discloses to the subsequent generation. By revealed preference, then, even the minimal disclosure equilibrium must yield a sufficiently high level of research productivity so that the payoffs received by each generation are preferable to a reversion to secrecy.[9] Overall, Proposition 4 highlights two features associated with Pasteur's Quadrant: (1) the benefits of establishing and sustaining an Open Science equilibrium when it is feasible and (2) the endogeneity of Open Science as an equilibrium when both Open Science and secrecy are feasible.

Proposition 4 compares the welfare associated with disclosure and secrecy equilibria. It is further possible to Pareto-rank different disclosure equilibria by the level of equilibrium research investment.

Proposition 5. When multiple stable path 'disclosure' equilibria exist, the researcher's payoffs are strictly increasing in the equilibrium level of research investment.

In a stable path disclosure equilibrium with optimal research investment x_1^1, the payoff of each generation of researcher is uniquely determined by size/quality of the knowledge pool, $z_1^1 = z(x_1^1)$, that each generation accesses. Moreover, each generation's payoff is increasing in z, since both the returns from citation and from private appropriation are increasing in z. A higher level of research investments on the equilibrium path implies a higher level of z accessed by each generation on the equilibrium path — the total payoffs for each generation are increasing in the equilibrium investment level.

Finally, it is straightforward that even the maximal level of research investment supported in an SPNE is (weakly) below the socially optimal level of research investment. In particular, the incentives for G_t to invest in research are limited to the private benefits received either through citations or commercial exploitation. However, G_t does not directly account for the value of its own investment in enhancing the research productivity of G_{t+1}. This positive intertemporal externality raises the optimal level of research investment for the social planner above that which can be supported by an SPNE.[10]

4. Extensions

Our analysis attempts to identify some of the key trade-offs associated with maintaining Open Science as an equilibrium over multiple research generations. In particular, the model focuses attention on the interdependence between the decision to build on knowledge from prior generations, the incentives for research investment, and the decision to disclose knowledge upon which subsequent generations might themselves build. To focus on these properties, we adopt several simplifying assumptions. We adopted a very specific form of knowledge accumulation process, and also limited our attention to a deterministic economic environment. Moreover, we have abstracted away from certain important institutional factors, such as the role played by intellectual property rights policy (e.g., patents over knowledge that also has scientific value). However, it is possible to enrich the current model to consider each of these concerns. While our current analysis of each of these issues is in no way comprehensive, we briefly review how the model might be enriched to account for several phenomena and policy concerns associated with the establishment and sustainability of Open Science over time.

4.1. Knowledge accumulation across multiple generations

The baseline model focuses on the case where knowledge produced in any one generation (and the choice of whether to disclose this information or not) has a direct impact only on the immediate next research generation. Scientific progress, of course, depends on the ability to draw on multiple prior generations, and the step-by-step process of knowledge accumulation is at the center of the

[9] Moreover, this welfare comparison may understate the welfare benefits arising from disclosure, because the welfare calculation only involves the utility accruing to researchers within the model. At least some of the "private" losses to G_t from disclosure (i.e., the discount in private benefits, λ) are likely associated with additional welfare gains to agents external to the model (such as the gains in consumer welfare arising from a more competitive product market for the innovations arising from discovery).

[10] As well, the social optimum will be associated with different values for λ^d,λ^{nd},F^d, and F^{nd}. However, the precise solution to the social planner's problem requires a detailed specification that includes a justification for the underlying objective function, a rational for the discount rate (as the optimal investment level will depend on the discount rate), etc. We leave a complete welfare analysis for future work.

modern theories of endogenous economic growth (e.g., Romer, 1990; Aghion and Howitt, 1992). Thus, one might be interested to know whether the key insights of our model continue to hold when the knowledge pool available to each generation of researchers reflects an accumulation of knowledge across multiple generations. Indeed, even in a more general environment where knowledge may accumulate over multiple generations, both disclosure and secrecy can be sustained as equilibria of the game depending on the underlying parameter values.

In order to extend our baseline model to the case of knowledge accumulation across multiple generations, two salient modeling issues need to be reconsidered: (1) strategic interaction across multiple generations, and (2) the exact nature of the accumulation of knowledge.

The issue of strategic interaction becomes more complex when knowledge is accumulated over multiple generations. The use of each generation's knowledge (and, therefore, its citation benefits) depends on the access decisions of each of the multiple future generations of researchers. To simplify the analysis, we assume that after a generation leaves the environment (i.e., at the end of the second period of his life) it cannot benefit from the citations offered by the future generations.[11]

Maintaining the assumption that the benefits from disclosing knowledge are limited to the benefits that one gets during one's own lifetime, it is useful to distinguish between two types of multi-generational knowledge accumulation processes. First, one may consider a completely general model where all prior generations impact the knowledge pool available to the current generation. Even if one assumes appropriate discounting, the treatment of this case is beyond the scope of this paper. It would require significant adjustments in the structure of the model, because, in such a setting, a steady-state *level* of knowledge pool (and, hence, a steady-state level of research investment) may not be a well-defined object, and one needs to solve for the steady-state growth *rates* of knowledge stock and research investments. Second, one can consider a finite period spillover, in which z_t depends on the investment and the disclosure decision of the past (finite) L generations as given below:

$$z_t = \sum_{l=1}^{L} d_{t-l} z(x_{t-l};\alpha). \tag{6}$$

The above formulation of z_t implies that size/quality of the knowledge pool available to G_t depends on the knowledge produced by each of the past L generations, $G_{t-1}, G_{t-2}, \ldots, G_{t-L}$, provided that they have decided to deposit the knowledge (i.e., $d_{t-1} = 1$). It is important to note that one can also consider a finite period spillover representing a so called "memoir effect." At the end of his "life", a researcher is indifferent between secrecy and disclosure since he cannot appropriate any additional returns. Thus, researchers may always choose to disclose their produced knowledge before they leave the environment (i.e., at the end of the second period of his life).[12] In such a case, we have

$$z_t = d_{t-1} z(x_{t-1};\alpha) + \sum_{l=2}^{L} z(x_{t-l};\alpha). \tag{7}$$

Observe that under this specification, even in a secrecy equilibrium the public knowledge pool is non-empty. As well, we keep all other aspects of the basic model (as described in Section 3) unchanged. It turns out that for both of these aforementioned knowledge accumulation processes (Eqs. (6) and (7)), some of the key insights developed in our basic model continue to hold. This issue is highlighted in the following proposition. (For the brevity of exposition, we only present the formal analysis under the formulation defined in Eq. (6). The proposition continues to hold with the formulation given in Eq. (7) while the proof requires some modifications.[13])

Proposition 6. When z_t is as given in Eq. (6), there exists values $F_*^{nd}(\lambda), F_*^d, \lambda_*^{nd}$ and $\lambda_*^d(F)$ such that a nondisclosure equilibrium necessarily exists for all (F,λ) if either $F \geq F_*^{nd}(\lambda)$ or $\lambda^{nd} \geq \lambda_*^{nd}$, and a disclosure equilibrium necessarily exists for $(F,\lambda) \leq (F_*^d, \lambda_*^d)$.

The above proposition suggests that both disclosure and secrecy can be supported as equilibria under suitable parameter regimes. Even if we consider a more general knowledge accumulation process, similar to our findings in the basic model, a disclosure equilibrium exists when *both* the access cost (F) and relative return from secrecy (λ) are sufficiently low. In contrast, a nondisclosure equilibrium exists when *either* λ *or* F is sufficiently high.

However, in contrast with the pervious analysis, Proposition 6 does not claim that a cyclic equilibrium must exist when a stable-path equilibrium fails to emerge. In fact, the dynamics of a cyclic equilibrium become more complicated (though qualitatively similar in spirit) compared to the basic model. This is due to the fact that the disclosure decision of the current generation affects not only the knowledge pool for the immediate next generation but also the pool available to the next L generations. This fact undermines analytical tractability of the basic model.[14]

4.2. Idiosyncratic payoffs

We first investigate the possibility that Open Science is subject to *hysteresis* — i.e., that it is more difficult to establish Open Science as equilibrium than to maintain that equilibrium once it is established. To do so, we allow the relative returns to secrecy (λ) to change over time. The private returns to scientific research may depend on the institutional and legal environment, and the strength of institutions or the legal regime may be different for different generations. In other words, we replace λ with a sequence of random variables λ_t where in each generation, λ_t is an independent draw from the probability distribution $g(\lambda)$ on $[1,\infty)$ with a finite mean. At the beginning of each generation, λ_t is realized, and the researcher chooses whether to access prior research, the level of investment in research, and

[11] This assumption, however, leaves room for the possibility that researchers in each generation may freely provide all knowledge at the end of their lives (say, a "memoir effect"), because they can no longer benefit from either the citation or the private exploitation. We will revisit this issue subsequently.

[12] In this context, it might be interesting to consider the role of a "market of ideas" where each generation of researcher has the option to sell her produced knowledge for a monetary return before she leaves the environment. However, a detailed analysis of the bargaining and monetary payments across generations is beyond the scope of this paper (see Gans et al. (2008) for a discussion of this issue.)

[13] Under the formulation given in Eq. (7), some modification is necessary for the proof for the existence of nondisclosure equilibrium. Because a researcher's output necessarily leaks out at the end of his life, even on a nondisclosure equilibrium, the knowledge pool is non-empty. This fact implies that the level of research investment in a nondisclosure equilibrium must also satisfy a 'fixed-point' property similar to the case of a disclosure equilibrium (see the proof of Proposition 1). However, the existence of such a fixed-point can be guaranteed by argument similar to the one used in the Proof of Proposition 1.

[14] In our basic model, a period of secrecy brings down the size/quality of the knowledge pool (z_t) for the next generation to zero. Thus, in order to establish a cyclic equilibrium, it is enough to show that starting from an empty knowledge pool, only a finite number of consecutive generations will choose to disclose. Once a single generation adopts secrecy, the knowledge pool becomes empty and the continuation game is identical to the game at the beginning of the cycle. Thus, the cycle continues to repeat itself giving rise to a "cyclic equilibrium." But, under the current formulation (as given by Eq. (6)), a period of secrecy does not pin down the value of z_t in the next period. In fact, the value of z_t continues to depend on the disclosure decision of all of the L prior generations. Therefore, every period of secrecy may not give rise to identical continuation games. Hence, an analytical proof of the existence of a cyclic equilibrium becomes significantly more challenging.

disclosure versus secrecy after observing λ_t (however, G_t cannot observe the realization of $\lambda_{t+k}, k=1,2,\ldots$).

Since the economic environment now changes from period to period, we can no longer confine the analysis to the class of stable path SPNE. Instead, we focus on the possibility that the equilibrium may exhibit hysteresis. To do so, we make the additional assumption that $\partial\mu^P/\partial z=0$, i.e., that the returns to private exploitation are independent of the quality/size of the knowledge pool (and, by implication, independent of whether that researcher has accessed the prior generation's knowledge pool).[15]

For a researcher with a given λ_t, the incentives for disclosure increase if the researcher in the prior generation have also disclosed. However, the disclosure choice of the prior generation depends on the value of λ_{t-1} (the returns to secrecy realized by the prior generation). When the benefits from citation are sensitive to z but the benefits from private exploitation are not, then the incentive to participate in Open Science for G_t will depend on whether there has been disclosure by G_{t-1} (which in turn depends on disclosure by generations prior to G_{t-1}). By implication, if the researcher from any one generation chooses *not* to disclose (e.g., they receive a very high λ_t), the incentives to disclose will be reduced in subsequent generations (and will remain so until a generation receives a sufficiently *low* value of λ to shift back to a disclosure regime). Denote $\Pr_{g(\lambda)}(d_t=1)$ as the probability that $d_t=1$ on the equilibrium path when λ follows the distribution $g(\lambda)$.

Proposition 7. Along any equilibrium path $\Pr_{g(\lambda)}(d_t=1|d_{t-1}=1)\geq\Pr_{g(\lambda)}(d_t=1|d_{t-1}=0)$.

In other words, the dynamics associated with idiosyncratic payoffs to secrecy result in *hysteresis* — for any equilibrium path, the probability of disclosing in generation t is higher if there has been disclosure in $t-1$, relative to the case where there has been secrecy in $t-1$.

4.3. The impact of patents

Finally, it is useful to consider the impact of formal IPR such as patents on the tradeoff between Open Science and Secrecy. At one level, patents combine elements from both regimes: patents require disclosure (and so share some of the attributes of Open Science by facilitating cumulative innovation (Scotchmer, 1991)), and provide protection from imitation (which could be interpreted as an increase in λ, the relative returns to Secrecy). However, a more careful analysis would incorporate patenting as a distinct strategy. In a follow-on paper, Gans et al. (2008) considers an environment in which, rather than simply modeling disclosure to the scientific literature versus complete secrecy, each generation faces a more complex set of disclosure options: secrecy (which involves neither patents nor publications), commercial science (where the only disclosures result from patenting), open science (where the only disclosures occur through scientific publication) and patent-paper pairs (where the firm discloses along both dimensions, as highlighted by Murray (2002)). While a complete discussion of that framework is beyond the scope of this paper, it is useful to highlight the key challenges which arise when attempting to incorporating patenting into a model that also allows for secrecy or scientific publication (and kudos from access and use by subsequent generations). First, the potential losses arising from disclosure through scientific publication will depend on whether G_t is also engaged in patenting and whether the information disclosure required for a patent is concordant (or not) with that required for scientific publication. Second, the motivation for scientific disclosure (and its relationship to patenting) can be better understood by disentangling the motives of researchers (who may place explicit value on participation in Open Science (Stern, 2004)) from the motives of investors of the research funds (who are primarily concerned about monetary returns). As discussed in Gans et al. (2008), patenting and publication are complementary to each other under particular conditions: when there is a high degree of overlap between the disclosures required for patenting and publication, and researchers are willing to incur a wage discount in order to publish in the scientific literature. Consequently, when patenting and publication are complements, increases in the strength of intellectual property may not simply enhance the returns to patenting but also facilitate scientific publication. An important potential implication of complementarity is that the design and effectiveness of intellectual property rights may have unintended consequences on the degree and impact of disclosures through Open Science.

Finally, an important difference between disclosures through publication and disclosures through patenting is that, under patenting, knowledge transfer across generations depends, at least in part, on a formal licensing agreement (Scotchmer, 1991, 2004; Gans et al., 2008). In our model, the access decision by G_{t+1} is an exogenous fixed cost, and is independent of the quality and level of disclosure by G_t. In other words, our findings regarding the conditions supporting Open Science are predicated on the idea that the returns to exploiting knowledge are increasing in the quality of that knowledge but the costs of acquiring knowledge are largely independent of quality. While this seems like a plausible assumption for scientific knowledge (indeed, the costs of learning may be declining when there is a particularly large advance that clarifies a particular research area), it is likely that, under most specifications of the bargaining environment, the cost of acquiring patented knowledge is increasing in the quality and importance of that knowledge. While a full treatment of the potential is beyond the scope of this paper, a licensing model would need to consider the nature of the bargaining between generations, and, in particular, how the bargaining between any two generations is likely to impact strategic interaction across subsequent generations (e.g., between G_{t+1} and G_{t+2}).

5. Conclusion

This paper is motivated by a simple yet important feature of cumulative knowledge production: to build upon prior discoveries, the knowledge underlying those discoveries must be disclosed and accessible. Whether knowledge gets disclosed to serve as an input into future research is not simply a function of the *type* of knowledge produced but depends on incentives. Researchers will endogenously choose whether to invest in prior knowledge, how much to invest in knowledge production, and whether to disclose knowledge for future use. More subtly, the incentives of any one researcher to disclose his own knowledge depends on the strategic choices of other researchers—indeed, it is possible that the selection of Open Science over Secrecy will depend on whether the researchers are able to coordinate on a favorable equilibrium outcome. More generally, the feasibility of Open Science is grounded in the microeconomic conditions under which research is conducted, and these conditions themselves depend on the public policy and the existing institutional and legal framework.

While the specific model we investigate yields several novel insights into the economic conditions supporting Open Science as an equilibrium, our analysis contains several key assumptions. Most importantly, our analysis is premised on the idea that scientists whose work is accessed by future generations are able to receive an exogenous benefit stream from that follow-on work when that work is used by future generations. We are explicitly agnostic about whether such rewards come in the form of prestige and stature (as

[15] Consistent with our prior discussion, this condition implies that, if λ were constant across time, a stable path SPNE will always exist (i.e., we can abstract away from k-period cyclic equilibria).

might be emphasized in the sociology of science (Merton, 1973)) or whether these rewards come in the form of income received through new employment opportunities, public expenditures on higher education, research grant funding, etc. (David, 1998; David and Keely, 2002; Lerner and Tirole, 2005). In particular, evaluating the impact of endogenous disclosure choices on labor markets is quite subtle (particularly when the information potentially revealed to the labor market need not be positive (as in, among others, Mukherjee (forthcoming)). The current analysis highlights the idea that, even with simple benefit function for follow-on research benefits, the feasibility of Open Science is grounded in the incentives for disclosure and strategic interaction among researchers. The analysis could be deepened by considering the endogeneity of citations (and prestige) to the institutional and strategic environment, with a particular focus on the subtle interdependency between the incentives for disclosures through Open Science and the incentives for disclosure through the patent system (as in Gans et al. (2008)).

A second important limitation of the analysis is that, by assuming that each generation can only build on knowledge from the immediately previous generation, we have mostly abstracted away from a range of important issues related to the accumulation of knowledge and access to knowledge over *multiple generations* (Rosenberg, 1982; Romer, 1990; Mokyr, 2002). As discussed in Section 4, we lose analytical tractability if we allow for complex strategic interactions among multiple generations. By focusing on a simple type of intertemporal linkage, we are able to precisely characterize the impact of the trade-off between secrecy and openness and the role of key parameters on this dynamic equilibrium process. The cumulative effect of strategic interaction when each generation may draw from *all* prior generations remains a question for future research.

Although our model abstracts away from a number of key aspects of Open Science, our analysis offers insight into the historical evidence regarding the rise of Open Science as an economic institution, and the benefits from Open Science for cumulative knowledge production in the process of economic growth (Romer, 1990; Mokyr, 2002). Our findings suggest that the near-continuous viability of scientific norms and scientific publication since the seventeenth century reflects more than simply public support for Science or the potential for breakthrough research. Instead, the ability to sustain a system with open exchange and disclosure of new discoveries has depended, among other things, on maintaining a sufficiently low cost of access to prior knowledge. Indeed, our analysis suggests that subsidies for specialized scientific education (e.g., postdoctoral training grants) may have a multiplier effect on maintaining Open Science. Subsidies not only reduce the private costs of accessing prior knowledge but also enhance incentives to disclose knowledge as well. Consistent with Jones (forthcoming), cumulative knowledge production depends as much on the ability to learn about prior discoveries as it does on a willingness to disclose one's own discoveries. At the same time, as a historically important system, the viability of science depends on maintaining an upper bound on the private financial returns that are achievable through secrecy. Ironically, the ability to maintain a system encouraging public disclosure requires that we place limits on the private exploitation of knowledge. Whether contemporary changes in research, such as the increased availability of intellectual property rights and increased corporate funding of basic projects, endanger Open Science remains an open yet fundamentally important question.

Appendix A

This appendix presents the proofs omitted in the text.

Proof of Proposition 1. We argue that there exist parameter values for which nondisclosure and disclosure equilibria exist. Moreover, if a stable path equilibrium does not exist, cyclic equilibrium must exist. The proof is given by the following steps.

Step 1. We first provide a sufficient condition for the existence of a nondisclosure equilibrium. We claim that $\{0,x_0^0,0\}_{t=-\infty}^{\infty}$ is an equilibrium if:[16]

$$U_t\left(0,x_0^0,0;0,0\right) \geq U_t\left(0,x_0^1,1;0,1\right). \tag{A1}$$

The argument is as follows. Consider the decision problem for G_t when $d_{t-1}=0$. Trivially, $\alpha_t=0$. Now $d_t=0$ if

$$U_t\left(0,x_0^0,0;0,0\right) \geq U_t\left(0,x_0^1,1;0,a_{t+1}\right) \quad \forall a_{t+1}.$$

But, by Assumption 2 (a), we get $U_t(0,x_0^1,1;0,1) \geq U_t(0,x_0^1,1;0,0)$. Hence Eq. (A1) provides a sufficient condition for 'nondisclosure' to be a best response to nondisclosure of a previous generation. As all generations face the same problem, when (A1) holds $\{0,x_0^0,0\}_{t=-\infty}^{\infty}$ is stable path equilibrium. Moreover this equilibrium is unique because x_0^0 is the unique argmax element. This is ensured by the concavity assumption on μ^P.

We now show that for suitable parameter values, a disclosure regime can also be supported as a stable path equilibrium.

Step 2. Let X be the value of x that solves $\bar{\mu}^C+\bar{\mu}^P=\psi x$. Since $\bar{\mu}^i$ is the upper bound on μ^i, $i=C,P$, without loss of generality, when $d_{t-1}=d_t=1$ and $a_t=a_{t+1}=1$, we can rewrite G_t's optimization problem as max $_x\mu^C(1,x_t,1;z_t,1)+\mu^P(1,x_t,1;z_t)-\psi x_t$ subject to $x_t\in[0,X]$. Since $z_t=z(x_{t-1})$, G_t's optimization problem can be rewritten as

$$\max_{x_t} f(x_t;x_{t-1}) \equiv \mu^C(1,x_t,1;z_t(x_{t-1}),1) + \mu^P(1,x_t,1;z_t(x_{t-1})) - \psi x_t.$$

Because μ^C and μ^P are concave, the solution to this optimization problem is unique for any given value of x_{t-1}. Let the solution be given by the function $\bar{x}(x_{t-1})$. As f is continuous in x_t and x_{t-1}, and the feasible set $[0,X]$ is compact, by the Theorem of Maximum, $\bar{x}$:$[0,X]\rightarrow[0,X]$ is continuous. So by Brouwer's Fixed Point Theorem we claim that there exists $x_1^1\in[0,X]$ such that $x_1^1=\bar{x}(x_1^1)$.

Step 3. Now, consider the candidate equilibrium $\{a_t=1,x_1^1,d_t=1\}$. To be an equilibrium, $\{a_t=1,x_1^1,d_t=1\}$ must solve Eq. (4) given $\{a_{t-1}=1,x_1^1,d_{t-1}=1\}$ and $a_{t+1}=1$. By construction, if G_{t-1} chooses x_1^1, G_t must also choose x_1^1. Other possible deviations are: $(a_t=0,d_t=1)$ and $d_t=0$ (i.e., $(a_t=0,d_t=0)$ and $(a_t=1,d_t=0)$). So it must be the case that

$$\begin{aligned}U_t\left(1,x_1^1,1;z_1^1,1\right) \geq \max\{&U_t\left(0,x_0^1,1,z_1^1,a_{t+1}\right),\\ &U_t\left(0,x_0^0,0,z_1^1,a_{t+1}\right),U_t\left(1,x_1^0,0,z_1^1,a_{t+1}\right)\}\end{aligned}$$

where $z_1^1=z(x_1^1)$. But note that $U_t(0,x_0^1,1,z_1^1,1) \geq U_t(0,x_0^1,1;z_1^1,0)$ (by Assumption 2(a)) and $U_t(a_t,x_{a_t}^0,0;z_1^1,a_{t+1}) \equiv U_t(a_t,x_{a_t}^0,0;z_1^1,0)$ (because the value of a_{t+1} is trivially 0 if $d_t=0$). Hence $\{a_t=1,x_1^1,d_t=1\}_{t=-\infty}^{\infty}$ is a stable path equilibrium sustaining disclosure regime if:

$$\begin{aligned}U_t\left(1,x_1^1,1;z_1^1,1\right) \geq \max\{&U_t\left(0,x_0^1,1;z_1^1,1\right),\\ U_t\left(0,x_0^0,0;z_1^1,0\right),&U_t\left(1,x_1^0,0;z_1^1,0\right)\}.\end{aligned} \tag{A2}$$

Finally, we argue that if neither nondisclosure nor disclosure regimes can be supported as equilibrium, a k-period cyclic equilibrium must exist.

Step 4. Note that the optimal choice of x_t can also be written as a function of z_t when $d_{t-1}=d_t=1$ and $a_t=a_{t+1}=1$. Let us denote this function as $x_t=h(z_t)$. Thus, for a given value of α we can write $z_t=z(x_{t-1};$

[16] Recall that we define the notation $x_a^d = \arg\max_x U_t(a_t,x,d_t;z_t,a_{t+1})$.

$\alpha)=z(h(z_{t-1});\alpha)$. Suppressing α for the sake of brevity, one can rewrite the above equation as $z_t=\bar{z}(z_{t-1})$. Moreover $\bar{z}$ is continuous since h and z are continuous. The rest of the proof relies on the continuity of $\bar{z}$.

Step 5. Because nondisclosure equilibrium does not exist, it must be the case that following $d_{t-1}=0$, G_t finds it optimal to choose $d_t=1$. That is,

$$U_t\left(0,x_0^0,0;0,0\right)<U_t\left(0,x_0^1,1;0,a_{t+1}\right).$$

But this is the case only if G_{t+1} finds it optimal to choose $a_{t+1}=1$ following $(a_t=0,d_t=1)$. Now, if it is also optimal for G_{t+1} to choose $d_{t+1}=0$, G_{t+2}'s decision problem is identical to that of G_t, because both generations face the same value of the available knowledge pool, i.e., $z_t=z_{t+2}=0$. Thus, an one-period cyclic equilibrium trivially exists, where only every alternate generation chooses to disclose.

But if it is optimal for G_{t+1} to choose $d_{t+1}=1$, we claim that following $(a_t=0,d_t=1)$, $\exists$ a finite $k\geq 2$ such that $(a_{t+\tau}=1, d_{t+\tau}=1)\forall\tau\in\{1,\ldots,k-1\}$ and $(a_{t+k}=1,d_{t+k}=0)$. Note that if for some k, $d_{t+k}=0$, the equilibrium action of G_{t+k+1} is the same as G_t. Thus, a k-period cyclic equilibrium exists as long as such a finite k exists.

Step 6. First, note that for any τ, if $d_\tau=1$ is a part of the best response of G_τ, it must be the case that $a_{\tau+1}=1$ is a part of the best response for $G_{\tau+1}$ (otherwise, G_τ is better off by resorting to secrecy and appropriating the payoff from private expropriation). Thus, if a cyclic equilibrium does *not* exist, it must be the case that following $(a_t=0, d_t=1)$, $G_{t+\tau}$ chooses $(a_{t+\tau}=1,d_{t+\tau}=1)$, $\forall\tau\in\{1,2,\ldots\}$. The next step shows that if disclosure equilibrium does not exist, such a sequence cannot be optimal for every generation.

Step 7. Fix a candidate equilibrium path for generations $\{G_{t+\tau}\}_{\tau=0}^{\infty}$ such that $(\alpha_t=0,d_t=1)$, and $G_{t+\tau}$ chooses $(\alpha_{t+\tau}=1,d_{t+\tau}=1)$, $\forall\tau\in\{1,2,\ldots\}$. We claim that the associated $\{x_{t+\tau}\}_{\tau=0}^{\infty}$ sequence is weakly increasing (we will give the proof shortly). Hence, $\{x_{t+\tau}\}_{\tau=0}^{\infty}$ must converge to a limit point, say, x^*. This follows from the fact that the optimal value of x_t is bounded in $[0,X]$. As z is an increasing function in x (by Assumption 5(b)), the associated $\{z_{t+\tau}\}_{\tau=0}^{\infty}$ must converge to a limit point, say, z^*. In Step 4 we argued that there exists a continuous function $\bar{z}$ such that $z_t=\bar{z}(z_{t-1})$. Thus, we must have $z^*=\bar{z}(z^*)$, i.e., z^* must be a fixed point for the mapping $\bar{z}$. But then, a disclosure equilibrium must exist where $(a_t=1,x_t=x^*,d_t=1)\forall t$. This contradicts our initial hypothesis that disclosure equilibrium does not exist. Finally, it remains to show that sequence $\{x_{t+\tau}\}_{\tau=0}^{\infty}$ is weakly increasing.

Step 8. We prove this claim by induction. First note that $z_t=0$. Let optimal $x_t=x_t^*$ where $(a_t=0,d_t=1)$. Now, $z_{t+1}>z_t=0$, and therefore the optimal $x_{t+1}=x_{t+1}^*>x_t^*$ when $(a_{t+1}=1,d_{t+1}=1)$ (by complementarity Assumption 4(d)). We now show that for an arbitrary generation $G_{t+\tau}$ such that $z_{t+\tau}>z_{t+\tau-1}$ and $(a_{t+\tau}=1,d_{t+\tau}=1)$ it must be the case that $x_{t+\tau}^*>x_{t+\tau-1}^*$. But this follows from the fact that the optimal value of x_t is an increasing function of z_t. □

Proof of Corollary 1. The proof is given in the following steps.

Step 1 (Threshold for λ). Observe that the only possible deviation from a nondisclosure path must involve $d_t=1$. Moreover, this deviation is profitable only if, following $d_t=1$, G_{t+1} chooses $a_{t+1}=1$ (this is due to the fact that $\lambda>1$). Condition (A1) ensures that even if $a_{t+1}=1$ following $d_t=1$, such deviation is not profitable. Note that (A1) can be written as

$$\max_{x_t}\left[\mu^P(0,x_t,0;0)-\psi x_t\right]=\max_{x_t}\left[\lambda\mu^P(0,x_t,1;0)-\psi x_t\right]\geq \max_{x_t}\left[\mu^C(0,x_t,1;0,1)+\mu^P(0,x_t,1;0)-\psi x_t\right]. \tag{A3}$$

The right-hand side of Eq. (A3) in increasing in λ. By continuity of $\lambda\mu^P-\psi x_t$, we claim that there exists a value of λ, say λ^{nd}, for which the above condition holds with equality. As this expression is increasing in λ, the above condition holds $\forall\lambda\geq\lambda^{nd}$, and, hence, nondisclosure is an equilibrium for all such λ.

Step 2 (Threshold for F). First, note that the optimal investment under secrecy, x_0^0, depends on λ. Now denote, $z_0^0(\lambda)=z(x_0^0(\lambda))$. Let $F_0^{nd}(\lambda)$ be the value of F (given λ) such that G_{t+1} is indifferent between accessing and not accessing the knowledge produced by G_t (while choosing x_t and d_t optimally) when G_{t+2} sets $a_{t+2}=0$ irrespective of G_{t+1}'s deposit decision. If G_{t+2} always sets $a_{t+2}=0$, G_{t+1} is clearly better off by not depositing the knowledge. Thus, $F_0^{nd}(\lambda)$ must solve the following equation:

$$\lambda\mu^P\left(0,x_0^0,1;z_0^0(\lambda)\right)-\psi x_0^0=\lambda\mu^P\left(1,x_1^0,1;z_0^0(\lambda)\right)-\psi x_1^0-F. \tag{A4}$$

Similarly, let $F_1^{nd}(\lambda)$ be the value of F (given λ) such that G_{t+1} is indifferent between accessing and not accessing the knowledge produced by G_t (while choosing x_t and d_t optimally) when G_{t+2} sets $\alpha_{t+2}=1$. Thus, $F_1^{nd}(\lambda)$ must solve the following equation:

$$\max_{x_t,d_t}U_{t+1}\left(0,x_{t+1},d_{t+1};z_0^0(\lambda),a_{t+2}=1\right)=\max_{x_t,d_t}U_{t+1}\left(1,x_{t+1},d_{t+1};z_0^0(\lambda),a_{t+2}=1\right),$$

that is,

$$\begin{aligned}&\max\{\mu^C\left(0,x_0^1,1;z_0^0(\lambda),1\right)+\mu^P\left(0,x_0^1,1;z_0^0(\lambda)\right)-\psi x_0^1,\\&\lambda\mu^P\left(0,x_0^0(\lambda),1;z_0^0(\lambda)\right)-\psi x_0^0(\lambda)\}\\&=\\&\max\{\mu^C\left(1,x_1^1,1;z_0^0(\lambda),1\right)+\mu^P\left(1,x_1^1,1;z_0^0(\lambda)\right)-\psi x_1^1,\\&\lambda\mu^P\left(1,x_1^0(\lambda),1;z_0^0(\lambda)\right)-\psi x_1^0(\lambda)\}-F.\end{aligned} \tag{A5}$$

Let $F^{nd}(\lambda)=\max\{F_0^{nd}(\lambda),F_1^{nd}(\lambda)\}$. So, for a given λ, at $F_i^{nd}(\lambda),i=0,1$, G_{t+1} is indifferent between drawing and not drawing from the knowledge pool given that G_t has invested x_0^0 amount in research, and G_{t+2} sets $a_{t+2}=i$. Thus, $\forall F\geq F^{nd}(\lambda), a_{t+1}=0$. Hence, for a given λ, secrecy is an equilibrium $\forall F\geq F^{nd}(\lambda)$ as there is no incentive for G_t to deposit. □

Proof of Corollary 2. The proof is given in the following steps.

Step 1 (Threshold for λ). Recall that Z^F is the set of all fixed points of the mapping $z_t=\bar{z}(z_{t-1})$. Consider a candidate disclosure equilibrium where the associated $z_t=z\in Z^F$. For any $F\leq F^d$, Eq. (A2) is satisfied if both of the following conditions hold:

$$\begin{aligned}&\max_{x_t}\left[\mu^C(1,x_t,1;z,1)+\mu^P(1,x_t,1;z)-\psi x_t\right]-F\\&\geq\max_{x_t}\left[\mu^P(0,x_t,0;z)-\psi x_t\right]\\&=\max_{x_t}\left[\lambda\mu^P(0,x_t,1;z)-\psi x_t\right],\end{aligned} \tag{A6}$$

and

$$\begin{aligned}&\max_{x_t}\left[\mu^C(1,x_t,1;z,1)+\mu^P(1,x_t,1;z,1)-\psi x_t\right]\\&\geq\max_{x_t}\left[\mu^P(1,x_t,0;z)-\psi x_t\right]\\&=\max_{x_t}\left[\lambda\mu^P(1,x_t,1;z)-\psi x_t\right].\end{aligned} \tag{A7}$$

The right-hand side of Eq. (A6) is increasing in λ. By continuity of μ^P, from Eq. (A6) we claim that given z, there exists a value of λ depending on F, say $\lambda^*(F,z)$, for which condition (A6) holds with equality. As this expression is increasing in λ, the above condition holds for $\forall\lambda\leq\lambda^*(F,z)$. Similarly, from Eq. (A7), we claim that there exists a value of λ, say $\bar{\lambda}(z)$, for which condition (A7) holds with equality and the inequality holds for $\forall\lambda\leq\bar{\lambda}(z)$.

Step 2 (Threshold for F). Take any $z \in Z^F$. From Eq. (A2) we know that if z supports a disclosure equilibrium, we must have $\max_{x_t} U_t(1,x_t,1;z,1) \geq \max_{x_t} U_t(0,x_t,1;z,1)$. This condition can be rewritten as:

$$\begin{aligned}&\max_{x_t}\left[\mu^C(1,x_t,1;z,1) + \mu^P(1,x_t,1;z) - \psi x_t\right] - \\ &\max_{x_t}\left[\mu^C(0,x_t,1;z,1) + \mu^P(0,x_t,1;z) - \psi x_t\right] \geq F.\end{aligned} \tag{A8}$$

Let $F(z)$ be the value of F for which the above inequality holds with equality. Define $F^d = \inf_{z \in Z^F} F(z)$. Now for all $F \leq F^d$ there exist a fixed point of the mapping $z_t = \bar{z}(z_{t-1})$ for which $(a_t = 0, d_t = 1)$ is not a profitable deviation.

Step 3 (Combining the thresholds). For all $F \leq F^d$, define $\lambda^d(F) = \sup_{z \in Z^F}\{\min\{\bar{\lambda}(z), \lambda^*(F,z)\}\}$ Hence for at least one $z \in Z^F$, Eqs. (A6), (A7), and (A8) are satisfied for any $F \leq F^d$ and $\lambda \leq \lambda^d(F)$. Therefore, Eq. (A2) holds $\forall (F,\lambda) \leq (F^d, \lambda^d(F))$. □

Proof of Lemma 1. We have already proved that $\lambda^d(F) = \sup_{z \in Z^F}\{\min\{\bar{\lambda}(z), \lambda(F,z)\}\}$ Moreover, $\lambda^*(F,z)$ is decreasing in F because the right-hand side of Eq. (A6) is increasing in λ. It remains to show that for any $z \in Z^F$ and $F \leq F^d$, $\lambda^{nd} \leq \lambda^*(F,z)$. Fix any $z = z^F \in Z^F$.

We will show that $\lambda^{nd} \leq \lambda^*(F^d, z^F)$. Recall that $\lambda^*(F^d, z)$ solves Eq. (A6) with equality, i.e.,

$$\underbrace{\max_{x_t}\left[\mu^C\left(1,x_t,1;z^F,1\right) + \mu^P\left(1,x_t,1;z^F\right) - \psi x_t\right]} = V - F^d = \underbrace{\max_{x_t}\left[\lambda\mu^P\left(0,x_t,1;z^F\right) - \psi x_t\right]}_{=W},$$

while λ^{nd} solves Eq. (A3) with equality, i.e.,

$$\underbrace{\max_{x_t}\left[\lambda\mu^P(0,x_t,1;0) - \psi x_t\right]}_{=W'} = \underbrace{\max_{x_t}\left[\mu^C(0,x_t,1;0,1) + \mu^P(0,x_t,1;0) - \psi x_t\right]}_{=V'}.$$

Note that when $a_t = 0$ then μ^P does not depend on the value of z_t. Hence $W' = W$ for all values of λ. Let $W' = W = W^{nd}$ for $\lambda = \lambda^{nd}$. By complementarity between a_t and x_t, we argue that $V \geq V'$. Moreover, from Eq. (A8), we know that $V - F^d \geq V'$. Therefore $V - F^d \geq V' = W^{nd}$. This is to say that Eq. (A6) is (strictly) satisfied when $\lambda = \lambda^{nd}$. Since W is increasing in λ, we must have $\lambda^*(F^d, z^F) \geq \lambda^{nd}$.

Now, it must be the case that $\lambda^{nd} \leq \lambda^*(F,z)$ for all $F \leq F^d$, because $\lambda^*(F,z)$ is decreasing in F. □

Proof of Proposition 2. The proof is given by the following steps:

Step 1. The optimal x under secrecy, x_0^0 must solve the following first-order condition:

$$\frac{\partial}{\partial x_t}\lambda\mu^P\left(0,x_0^0,1;z_t\right) \equiv \psi.$$

Because μ^P is concave in x, an increase in ψ decreases x_0^0. Next, note that α affects the first-order condition only through its impact on z_t. Because $a_t = 0$ for all t, z_t does not affect $\partial\mu^P/\partial x_t$ (by Assumption 4 (d)). Thus x_0^0 is independent of α.

Step 2. The optimal x under disclosure, x_1^1 must solve the following first-order condition:

$$\frac{\partial}{\partial x_t}\mu^C\left(1,x_1^1,1;z\left(x_1^1\right),1\right) + \frac{\partial}{\partial x_t}\mu^P\left(1,x_1^1,1;z\left(x_1^1\right)\right) \equiv \psi.$$

Taking the total derivative with respect to ψ, one arrives at:

$$\left[\mu_{xx}^C + \mu_{xx}^P + \left(\mu_{xz}^C + \mu_{xz}^P\right)\frac{\partial z}{\partial x}\right]\frac{\partial x_1^1}{\partial \psi} = 1.$$

Step 3. Now, in any dynamically stable equilibrium, $[\mu_{xx}^C + \mu_{xx}^P + (\mu_{xz}^C + \mu_{xz}^P)\partial z/\partial x] < 0$. To see this, recall that the optimal x_t as a function of x_{t-1} is given by the continuous function $\bar{x}(x_{t-1})$, as defined in the Step 2 of the proof of Proposition 1. If the x_1^1 can be supported as a dynamically stable equilibrium, it must be the case that $x_1^1 = \bar{x}(x_1^1)$ and $\partial\bar{x}/\partial x_{t-1} < 1$ at $x_{t-1} = x_1^1$ (i.e., the mapping x must intersect the 45° line from above at the fixed point). But from the first order condition that yields $\bar{x}$, one obtains that $(\mu_{xx}^C + \mu_{xx}^P)\partial\bar{x}/\partial x_{t-1} + (\mu_{xz}^C + \mu_{xz}^P)\partial z/\partial x_{t-1} = 0$. Now, $\mu_{xx}^C + \mu_{xx}^P < 0$. Therefore, if $\partial\bar{x}/\partial x_{t-1} < 1$, we must have $[\mu_{xx}^C + \mu_{xx}^P + (\mu_{xz}^C + \mu_{xz}^P)\partial z/\partial x] < 0$. Hence, $\partial x_1^1/\partial\psi < 0$.

Step 4. Finally, taking the total derivative of the first-order condition with respect to α, one arrives at

$$\left[\mu_{xx}^C + \mu_{xx}^P + \left(\mu_{xz}^C + \mu_{xz}^P\right)\frac{\partial z}{\partial x}\right]\frac{\partial x_1^1}{\partial \alpha} + \left(\mu_{xz}^C + \mu_{xz}^P\right)\frac{\partial z}{\partial \alpha} = 0.$$

Now, we have already argued that $[\mu_{xx}^C + \mu_{xx}^P + (\mu_{xz}^C + \mu_{xz}^P)\partial z/\partial x] < 0$, and $(\mu_{xz}^C + \mu_{xz}^P)\partial z/\partial\alpha > 0$ by Assumptions 4(d) and 5(b). Thus, $\partial x_1^1/\partial\alpha > 0$. □

Proof of Proposition 3a. Recall that $\lambda^d, \lambda^{nd}, F^d$, and F^{nd} are determined from the following equations:

a. $F^{nd} = \max\{F_0^{nd}, F_1^{nd}\}$, where F_0^{nd} solves Eq. (A4) and F_1^{nd} solves Eq. (A5).
b. λ^{nd} solves Eq. (A3) with equality.
c. $F^d = \inf_{z \in Z^F} F(z)$, where $F(z)$ is the value of F for which the Eq. (A8) holds with equality.
d. $\lambda^d(F) = \sup_{z \in Z^F}\{\min\{\bar{\lambda}(z), \lambda^*(F,z)\}\}$, where $\bar{\lambda}$ solves (A7) with equality and $\lambda^*(F,z)$ solves (A6) with equality.

The rest of the proof is given in the following steps.

Step 1 (Impact on F^{nd}). To see the impact of α on F^{nd}, observe that α enters only in the right-hand side of the Eqs. (A4) and (A5) through z_0^0 (i.e., the value of z_{t+1}) (recall that by Assumption 4(a), the left-hand sides of Eqs. (A4) and (A5) are both independent of z_{t+1} because $a_{t+1} = 0$). Moreover the right-hand sides of the Eqs. (A4) and (A5) are increasing in α. Therefore, to maintain these equality constraints, both F_0^{nd} and F_1^{nd} must increase. Hence, F^{nd} increases as well.

Step 2 (Impact on λ^{nd}). As α does not affect Eq. (A3), λ^{nd} is independent of α.

Step 3 (Impact on F^d). The left-hand side of Eq. (A8) must increase with α. To see this, consider the total derivative of the left-hand side of Eq. (A8) with respect to α. Using the Envelope Theorem and the fact that both μ^C and μ^P are independent of z (and hence, of α) if $a_t = 0$, the total derivative of the left-hand side of Eq. (A8) with respect to α can be written as:

$$\left[\mu_z^C\left(1,x_1^1,1;z,1\right) + \mu_z^P\left(1,x_1^1,1;z\right)\right]\left[\frac{\partial z}{\partial x}\cdot\frac{\partial x_1^1}{\partial \alpha} + \frac{\partial z}{\partial \alpha}\right].$$

This expression is positive because each term is individually positive. Thus, F^d must increase with α.

Step 4 (impact on λ^d). Both the left- and the right-hand side of Eq. (A7) are increasing in α. However, one cannot rank which side increases more, and hence the impact of α on $\bar{\lambda}$ (and hence, on λ^d) is ambiguous. However, if $\partial\mu^P/\partial z$ is sufficiently small compared to $\partial\mu^P/\partial z$, then the left-hand side of Eq. (A7) increases more than the right-hand side. So, $\bar{\lambda}$ must increase. Also note that $\lambda^*(F,z)$ is increasing in α, because the right-hand side of Eq. (A6) is increasing in α, but the left-hand side is independent of α (because $a_t = 0$). Thus, λ^d must increase in α if $\partial\mu^P/\partial z$ is sufficiently small compared to $\partial\mu^P/\partial z$. □

Proof of Proposition 3b. We prove this result by showing that the comparative statics of each of the four boundary points involves a comparison of the optimal research investment, x_a^d, under different disclosure and access regimes and that such comparisons are often ambiguous.

Step 1 (Impact on F^{nd}). Consider the Eq. (A5) that defines F_1^{nd}. Suppose, that the underlying parameter values are such that (A5) $\mu^C(0,x_0^1,1;z_0^0,1) + \mu^P(0,x_0^1,1;z_0^0) - \psi x_0^1 > \lambda\mu^P(0,x_0^0,1;z_0^0) - \psi x_0^0$, but $\mu^C(1,x_1^1,1;z_0^0,1) +$

$\mu^P(1,x_1^1,1;z_0^0)-\psi x_1^1<\lambda\mu^P(1,x_1^0,1;z_0^0)-\psi x_1^0$. Thus, Eq. (A5) boils down to

$$F_1^{nd}=\mu^C\left(0,x_0^1,1;z_0^0,1\right)+\mu^P\left(0,x_0^1,1;z_0^0\right)-\psi x_0^1-\lambda\mu^P\left(1,x_1^0,1;z_0^0\right)+\psi x_1^0.$$

Now,

$$\frac{\partial F_1^{nd}}{\partial\psi}=-\mu_z^P\left(1,x_1^0,0;.\right)\frac{\partial z}{\partial x}\frac{\partial x_1^0}{\partial\psi}-\left(x_0^1-x_1^0\right).$$

But $\partial F_1^{nd}/\partial\psi$ cannot be signed because $(x_0^1-x_1^0)$ cannot be signed. Hence, the impact of ψ on F^{nd} is also ambiguous.

Step 2 (Impact on λ^{nd}). Taking total derivative of both sides of Eq. (A3) with respect to ψ one arrives at

$$\mu^P\left(0,x_0^0,1;0\right)\frac{\partial\lambda^{nd}}{\partial\psi}-x_0^0=-x_0^1$$

Again, $\partial\lambda^{nd}/\partial\psi$ cannot be signed as $(x_0^1-x_1^0)$ cannot be signed.

Step 3 (Impact on F^d). Taking total derivative of both sides of Eq. (A8) with respect to ψ one arrives at

$$\frac{\partial F^d}{\partial\psi}=\left[\mu_x^C+\mu_x^P\right]\frac{\partial z}{\partial x}\frac{\partial x_1^1}{\partial\psi}-\left(x_1^1-x_0^1\right).$$

Because $\partial x_1^1/\partial\psi<0$ and $(x_0^1-x_1^0)>0$, $\partial F^d/\partial\psi<0$.

Step 4 (Impact on λ^d). Consider the Eq. (A7) that defines $\bar{\lambda}(z)$. Taking total derivative of both sides of Eq. (A7) with respect to ψ, one arrives at

$$\left[\mu_z^C\left(1,x_1^1,1;.\right)+\mu_z^P\left(1,x_1^1,1;.\right)\right]\frac{\partial z}{\partial x}\frac{\partial x_1^1}{\partial\psi}-x_1^1=\mu^P\left(1,x_1^0,1;.\right)\frac{\partial\bar{\lambda}}{\partial\psi}+\mu_z^P\left(1,x_1^0,1;.\right)\frac{\partial z}{\partial x}\frac{\partial x_1^0}{\partial\psi}-x_1^0.$$

Now, $\partial\bar{\lambda}/\partial\psi$ cannot be signed as, among other terms in the above expression, $(x_1^1-x_1^0)$ cannot be signed. □

Proof of Proposition 4. Let z be the level of z_t associated with a given disclosure equilibrium. Existence of a disclosure equilibrium requires $U_t(1,x_1^1,1;z_t,1)\geq U_t(0,x_0^0,0;z_t,0)$. As $\alpha_t=0$, $\max_{x_t}U_t(0,x_t,0;z_t,0)$ is independent of z_t. Hence, $U_t(0,x_0^0,0;z_t,0)\equiv U_t(0,x_0^0,0;0,0)$. Therefore, when a disclosure equilibrium exists, we must have $U_t(1,x_1^1,1;z,1)\geq U(0,x_0^0,0;0,0)$. □

Proof of Proposition 5. Consider a stable path disclosure equilibrium with an associated $z=z^*$. Let the payoff to G_t from such equilibrium be $U_t(z_1)$, where

$$U_t\left(z^*\right)=\max_{x_t}\mu^C\left(1,x_t,1;z^*,1\right)+\mu^P\left(1,x_t,1;z^*\right).$$

Consider another stable path disclosure equilibrium where the associated $z=z^{**}>z^*$. Now $U_t(z^{**})>U_t(z^*)$ as, by Envelop Theorem, $U_t^{'}(z)=\mu_z^C+\mu_z^P>0$ (by Assumption 4 (b)). □

Proof of Proposition 6. This proof closely follows the proof of Proposition 1 and its corollaries. Therefore, for the sake of brevity, we will only elaborate on the additional nuances that are introduced by the multigenerational knowledge accumulation process.

Step 1. First, consider the proof of the existence of a 'nondisclosure' equilibrium. Note that even under Eq. (6), $z_t=0$ when $d_{t-1}=0$. Therefore, the proof of the existence of a 'nondisclosure equilibrium' is identical to the proof given in the Step1 of Proposition 1. Also, the threshold parameter values, $F_*^{nd}(\lambda)$ and λ_*^{nd}, can be derived as given in the proof of Corollary 1.

Step 2. Next, we prove the existence of a disclosure equilibrium. This proof is based on a modification of Step 2 and Step 3 of the proof of Proposition 1. The modification is necessary because under Eq. (6), in any candidate disclosure equilibrium, z_t depends on the vector of past research investments $(x_{t-1},x_{t-1},\ldots,x_{t-L})$ rather than only the last generation's investment x_{t-1}. Thus, following the analysis is Step 2 of Proposition 1, G_t's optimization problem under Eq. (6) can be rewritten as

$$\max_{x_t}f(x_t;x_{t-1},\ldots,x_{t-L})\equiv\mu^C(1,x_t,1;z(x_{t-1},\ldots,x_{t-L}),1)+\mu^P(1,x_t,1;z(x_{t-1},\ldots,x_{t-L}))-\psi x_t.$$

Step 3. Define $h(x;y)=f(x;y,\ldots,y)$. When $(x_{t-1},x_{t-2},\ldots,x_{t-L})=(x^*,x^*,\ldots x^*)$, i.e., all of the L prior generations invest x^*, G_t's optimization problem can be written as

$$\max_{x_t}h\left(x_t;x^*\right)\equiv\mu^C\left(1,x_t,1;z\left(x^*,\ldots,x^*\right),1\right)+\mu^P\left(1,x_t,1;z\left(x^*,\ldots,x^*\right)\right)-\psi x_t$$

Now, we can apply Brouwer's Fixed Point Theorem as we did in Step 2 of the proof of Proposition 1. Because μ^C and μ^P are concave, the solution to this optimization problem is unique for any given x^*. Let the solution be given by the function $\bar{x}(x^*)$. As h is continuous in x_t and x^*, and the feasible set $[0,X]$ is compact, by the Theorem of Maximum, $\bar{x}:[0,X]\to[0,X]$ is continuous. So by Brouwer's Fixed Point Theorem, we claim that there exists $x_1^1\in[0,X]$ such that $x_1^1=\bar{x}(x_1^1)$.

Step 4. Now, consider the candidate equilibrium $\{a_t=1,x_1^1,d_t=1\}$. To be an equilibrium, $\{a_t=1,x_1^1,d_t=1\}$ must solve (4) given $\{a_{t-1}=1,x_1^1,d_{t-1}=1\}$ and $a_{t+1}=1$. By construction, if $G_{t-1},\ldots,G_{t-L}$ all choose x_1^1, G_t must also choose x_1^1. Other possible deviations are: $(a_t=0,d_t=1)$ and $d_t=0$ (i.e., $(a_t=0,d_t=0)$ and $(a_t=1,d_t=0)$). So it must be the case that

$$U_t\left(1,x_1^1,1;z_1^1,1\right)\geq max\{U_t\left(0,x_0^1,1,z_1^1,a_{t+1}\right),U_t\left(0,x_0^0,0,z_1^1,a_{t+1}\right),U_t\left(1,x_1^0,0,z_1^1,a_{t+1}\right)\}$$

where $z_1^1=z(x_1^1,\ldots x_1^1)$. The above 'no deviation' condition yields the threshold values F_*^d and $\lambda_*^d(F)$ following the arguments presented in the proof of Corollary 2. □

Proof of Proposition 7. The proof is given in the following steps.

Step 1. Fix a value of F and an equilibrium strategy sequence $\{\sigma_t\}_{t=-\infty}^{\infty}$. Consider a generation G_{t-1} that has not disclosed, and hence $z_t=0$. Given the strategy of G_{t+1}, σ_{t+1}, the optimization problem for G_t is:

$$\max_{x_t,d_t}E\left[\mu^C(0,x_t,d_t;0,\bullet)+\mu^P(0,x_t,d_t;0)\right]-\psi x_t,$$

where the expectation of μ^C is taken over the probability distribution on a_{t+1} induced by σ_{t+1}. We can rewrite this optimization problem in two separate optimization problems, each associated with a different value of d_t, and G_t chooses the (x_t,d_t) tuple that maximizes his expected payoff. So for $d_t=1$, the problem boils down to

$$\max_{x_t}E\left[\mu^C(0,x_t,1;0,\bullet)+\mu^P(0,x_t,1;0)\right]-\psi x_t,$$

where as, for $d_t=0$, the problem is

$$\max_{x_t}\lambda\mu^P(0,x_t,1;0)-\psi x_t.$$

Step 2. Let the payoff to G_t associated with each of these two optimization problems be $U^1(z=0)$ and $U^0(z=0,\lambda)$. As U^0 is increasing in λ and is strictly less than $U^1(z=0)$ for $\lambda=1$, there exists a value of λ, say λ_0, such that $U^1(z=0)=U^0(z=0,\lambda_0)$. Hence, G_t chooses $d_t=1$ for all $\lambda\geq\lambda_0$. Therefore, $\Pr_{g(\lambda)}(d_t=1|d_{t-1}=0)=\Pr(\lambda<\lambda_0)$.

Step 3. Now, suppose, G_{t-1} discloses his knowledge output. So z_t rises from 0 to a positive quantity, say z^*. If F is high enough to ensure that G_t will set $a_t=0$, the G_t's optimization problem remains unchanged from above. Therefore, the result trivially holds. Suppose, with disclosure by G_{t-1}, G_t finds it optimal to set $a_t=1$. But then, $E\mu^C(1,x_1^1,1;z^*,\bullet)>E\mu^C(1,x_1^1,1;0,\bullet)$ in any equilibrium. The argument lies in the fact that x_t and a_t are complements in G_t's objective function, and the payoff of G_{t+1} from accessing G_t's knowledge output is increasing in x_t and α_t. Thus, fixing F, if for some λ, $a_{t+1}=1$ under σ_{t+1} and $z_t=0$, then σ_{t+1} must also induce $a_{t+1}=1$ when $z_t=z^*>0$.

Step 4. Let λ^* be the value of λ for which $U^1(z=z^*)=U^0(z=z^*,\lambda^*)$. Therefore, $\Pr_{g(\lambda)}(d_t=1|d_{t-1}=1)=\Pr(\lambda<\lambda^*)$. As $\partial\mu^P/\partial z_t=0$, and $E\mu^C(1,x_t,1;z^*,\bullet)>E\mu^C(1,x_t,1;0,\bullet)$, it must be the case that $\lambda^*>\lambda_0$. Thus, $\Pr(\lambda<\lambda^*)>\Pr(\lambda<\lambda_0)$. □

References

Anton, J., Yao, D., 2004. Little patents and big secrets: managing intellectual property. RAND Journal of Economics 35 (1), 1–22.
Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60 (2), 323–351.
Aghion, P., Dewatripont, M., Stein, J., 2005. Academic Freedom, Private-Sector Focus, and the Process of Innovation. NBER Working Paper No. W11542.
Arora, A., 1995. Licensing tacit knowledge: intellectual property rights and the market for know-how. Economics of Innovation and New Technology 4, 41–49.
Arrow, K., 1962. Economics of welfare and the allocation of resources for invention. In: Nelson, R.R. (Ed.), The Rate and Direction of Inventive Activity. Princeton University Press, Princeton.
Blumenthal, D., Campbell, E.G., Anderson, M.S., Causino, N.A., Louis, K.S., 1997. Withholding research results in academic life science: evidence from a National Survey of Faculty. Journal of the American Medical Association 227 (16), 1224–1228.
Dasgupta, P., David, P., 1994. Towards a new economics of science. Research Policy 23 (5), 487–521.
David, P., 1998. Common agency contracting and the emergence of Open Science institutions. AEA Papers and Proceedings 88 (2), 15–21.
David, P., 2004. Patronage, reputation, and common agency contracting in the scientific revolution: from keeping 'nature's secrets' to the institutionalization of 'Open Science'. SIEPR Policy Paper No. 03-039.
David, P., Keely, L., 2002. The economics of scientific research coalitions: collaborative network formation in the presence of multiple funding agencies. SIEPR Discussion Paper No. 01-24.
Furman, J., Stern, S., 2008. Climbing atop the shoulders of giants: the impact of institutions on cumulative research. NBER Working Paper No. 12523.
Gans, J., Murray, F., and Stern, S., (2008). "Patents, Papers & Privacy: The disclosure of scientific and commercial knowledge," mimeo Melbourne Business School.
Gallini, N., Scotchmer, S., 2002. In: Jaffe, Adam, Lerner, Josh, Stern, Scott (Eds.), Intellectual Property: When is it the Best Incentive Mechanism? Innovation Policy and the Economy, Vol 2.
Geison, G.L., 1995. The Private Science of Louis Pasteur. Princeton University Press, Princeton, NJ.
Gossner, O., 1996. Overlapping generation games with mixed strategies. Mathematics of Operation Research Vol. 21 (2), 477–486.
Grossman, G., Helpman, E., 1991. Innovation and Growth in the Global Economy. MIT Press, Cambridge.
Horstman, I., MacDonald, G., Slivinski, A., 1985. Patents as information transfer mechanisms: to patent or (maybe) not to patent. Journal of Political Economy 93 (5), 837–858.
Jones, C., 1995. R&D based models of economic growth. Journal of Political Economy 103 (4), 739–784.
Jones, B. (2008). "The Burden of Knowledge and the Death of the Renaissance Man: Is Innovation Getting Harder?" Review of Economic Studies (forthcoming).
Kevles, D., 2002. Of Mice and Money: The Story of the World's First Animal Patent. Daedalus. spring 2002.
Kultti, K., Takaloz, T., Toikkax, J., 2007. Secrecy versus patenting. RAND Journal of Economics 38 (1), 22–42.
Lacetera, N. (2008). Academic Entrepreneurship, mimeo, Weatherhead School of Management.
Lerner, J., Tirole, J., 2005. The economics of technology sharing: open source and beyond. Journal of Economic Perspectives 19, 99–120.
Merton, R., 1973. The Sociology of Science: Theoretical and Empirical Investigation. University of Chicago Press, Chicago.
Mokyr, J., 2002. The Gifts of Athena: The Historical Origins of the Knowledge Economy. Princeton University Press, Princeton.
Mowery, D., Rosenberg, N., 1998. Paths of Innovation: Technological Change in 20th Century America. Cambridge University Press, Cambridge, UK; New York.
Mukherjee, A. (2008). "Career Concerns, Matching, and Optimal Disclosure Policy," International Economic Review (forthcoming).
Murray, F., 2002. Innovation as overlapping scientific and technological trajectories: exploring tissue engineering. Research Policy 31 (8-9), 1389–1403.
Murray, F. (2006). "The Oncomouse that Roared: Resistance & Accommodation to Patenting in Academic Science," mimeo, MIT.
Murray, F., Stern, S., 2007. Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis. Journal of Economic Behavior and Organization 63 (4), 648–687.
Nelson, R., 1959. The simple economics of basic scientific research. Journal of Political Economy 67 (2), 297–306.
Nelson, R., Rosenberg, N., 1994. American universities and technical advance in industry. Research Policy 23 (3), 323–348.
Polanyi, M., 1967. The Tacit Dimension. Anchor Books, New York.
Romer, P., 1990. Endogenous technological change. Journal of Political Economy 98 (5), S71–S102.
Rosenberg, N., 1974. Science, invention, and economic growth. Economic Journal 84 (333), 90–108.
Rosenberg, N., 1982. Inside the Black Box. Cambridge University Press, Cambridge.
Rosenberg, N., 1990. Why do firms do basic research (with their own money). Research Policy 19 (2), 165–174.
Scotchmer, S., 1991. Standing on the shoulders of giants: cumulative research and the patent law. Journal of Economic Perspectives 5 (1), 29–41.
Scotchmer, S., 2004. The political economy of intellectual property treaties. Journal of Law, Economics and Organizations 20 (2), 415–437.
Stern, S., 2004. Do scientists pay to be scientists? Management Science 50 (1), 835–853.
Stokes, D., 1997. Pasteur's Quadrant: Basic Science and Technological Innovation. The Brookings Institution, Washington D.C.