

Informational externalities and strategic interaction

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von

MIRJAM SALISH

aus Wien

Bonn, 2016

Dekan: Prof. Dr. Daniel Zimmer
Erstreferent: Prof. Sven Rady, PhD
Zweitreferent: Prof. Dr. Dennis Gärtner

Tag der mündlichen Prüfung: 18.11.2016

Diese Dissertation ist auf dem Hochschulschriftenserver der ULB Bonn http://hss.ulb.uni-bonn.de/diss_online/ elektronisch publiziert.

Acknowledgement

I would like to thank my main advisor Prof. Sven Rady not only for many helpful discussions and continuous feedback, which have helped to considerably improve this thesis, but also for his encouragement and support. Further, I would like to use this opportunity to thank my second advisor Prof. Dennis Gärtner as well as Prof. Piero Gottardi, my advisor during a research visit at the European University Institute in Florence. This thesis has greatly benefited from their suggestions and feedback.

I am particularly grateful to the support from the Deutsche Forschungsgemeinschaft through SFB/TR 15. Being a member of the SFB/ TR 15 research group has not only provided the financial support necessary to write this thesis, but also given me the opportunity to present my work at conferences and receive outstanding feedback (particularly at the SFB/ TR 15 conferences) at different stages of my thesis.

Finally, I want to thank the Bonn Graduate School of Economics for providing an excellent research environment. Many members of the Bonn Graduate School of Economics have in one way or another helped me to write this thesis. Some through administrative assistance, some through content-related feedback (particularly in the research workshops) and many through numerous interesting and entertaining discussions related to research and beyond.

Contents

Introduction	1
1 Learning faster or more precisely? Strategic experimentation in networks	6
1.1 Introduction	6
1.2 Model	10
1.2.1 A two-armed bandit model	11
1.2.2 Introducing a network structure	12
1.2.3 Timing and information structure	13
1.2.4 Strategies and equilibrium concept	15
1.3 The empty and the complete network	16
1.3.1 The empty network	16
1.3.2 The complete network	18
1.4 The ring network	20
1.5 The star network	25
1.6 Welfare analysis	32
1.6.1 The optimal network	32
1.6.2 The complete network with a specialist	36
1.6.3 Numerical example	37
1.7 Discussion	39
2 Innovation, intersectoral R&D spillovers and intellectual property rights	42
2.1 Introduction	42
2.1.1 Related literature	46
2.1.2 Examples	48

2.2	Model	48
2.2.1	Timing	49
2.2.2	Intrasectoral R&D spillovers	50
2.2.3	Beliefs and informational externalities	51
2.2.4	Producers, consumers and welfare	52
2.2.5	Investment strategies and equilibrium concept	53
2.3	R&D investment with intersectoral R&D spillovers	53
2.3.1	One firm in each sector	54
2.3.2	Two firms in one sector	57
2.4	Selecting a sector	61
2.5	Intersectoral R&D spillovers and intellectual property rights	65
2.6	Discussion	66
3	The transparency of the patent system and its impact on innovation	68
3.1	Introduction	68
3.1.1	Related literature	72
3.2	Model	74
3.2.1	Investment in uncertain research projects	74
3.2.2	The value of a patent and R&D spillovers	75
3.2.3	Beliefs and informational spillovers	75
3.2.4	Information structure and timing	76
3.2.5	Strategies and equilibrium concept	79
3.2.6	Consumers	79
3.3	The monopolist	79
3.4	Research cooperation	81
3.5	Research competition	84
3.5.1	Winner-takes-all competition	84
3.5.2	Perfect R&D spillovers	91
3.5.3	Sequential innovation	98
3.6	Conclusion	102
	Appendix A	103

Appendix B	114
Appendix C	121
Bibliography	134

List of Figures

1.1	Network structures	13
1.2	The star network	14
1.3	Expected payoffs of experimenting for the single agent	18
1.4	Equilibrium experimentation in the complete network	20
1.5	Equilibrium experimentation in the ring network	24
1.6	Equilibrium experimentation in the star network	28
1.7	The ring and the star network in comparison	31
1.8	The optimal network	35
1.9	Optimal networks - numerical example	38
2.1	R&D investment with intersectoral spillovers	58
2.2	R&D investment in sector s	59

List of Tables

3.1	Payoff matrix for R&D investment in stage 2	98
-----	---	----

Introduction

Uncertainty is a feature of many economic decision problems, processes and situations of economic interaction. Through research and learning agents reduce this uncertainty and thereby create informational externalities to others. For example, consumers trying a new product (for example, a new movie or electronic device) or a novel service share their experiences with friends and acquaintances, and even with strangers through customer reviews in the internet. A recommendation by a friend, say for a movie, reduces the uncertainty about the quality of the movie, as this friend's opinion provides a signal about the movie's quality. In this way the experience of others influences our opinions and consequently our decisions. Another example in which informational externalities play an important role is research or innovation. Firms invest in R&D in the search of new technologies. However, inventions of one firm often reveal information or knowledge to other (possibly competing) firms that have been acquired by the inventor in a lengthy and expensive process. Without proper protection of its intellectual property firms may find their invention copied and possibly improved shortly after releasing the original innovation, because of the informational externalities created by their discovery.

Economic agents that interact strategically take into account the informational externalities created by themselves and others when deciding which products to buy or whether to invest in R&D. Models of strategic experimentation are a useful tool to study and analyze strategic interaction with uncertainty and informational externalities. The idea of these models is that agents can invest money or effort into a risky project (often referred to as the risky arm of a multi-armed bandit machine), where ex ante it is not clear whether the payoff from the risky alternative is higher than the payoff from investing into a safe alternative with known return. By trying out the risky alternative repeatedly (that is, by experimenting) agents learn about its quality from observing the outcomes of their

own actions, but also from observing the outcomes of others' actions. In a situation where agents can observe each others experimentation decisions and outcomes, a public good problem arises and agents have an incentive to free-ride on the experimentation effort of others.

The first models of strategic experimentation were introduced in Bolton and Harris (1999) and Keller, Rady and Cripps (2005). In the first, the payoff generating process follows a Brownian motion with drift. Additional to the free-riding effect an encouragement effect arises, which means that the experimentation of one agent motivates other agents to experiment. In Keller et al. (2005) agents face an exponential bandit machine in which the first high payoff realization reveals that the risky alternative is good. These two papers laid the foundation for a growing body of literature on strategic experimentation and particularly the exponential bandit framework has been modified and studied in different settings. A more detailed discussion of the related literature can be found in the corresponding chapters of the thesis.

The aim of this thesis is to contribute to our understanding of uncertainty, learning and informational externalities in strategic interaction. To be more precise, the following questions are answered: How are the incentives of strategic agents to invest in risky and innovative activities affected by the monitoring protocol and monitoring imperfections? In particular, under which conditions does a fast and perfect diffusion of information maximize welfare and when can monitoring imperfections or incomplete interaction structures be beneficial? Different monitoring structures can be represented in the form of social or economic networks. Comparing different network structures shows how experimentation effort is influenced by a delay in information transmission or who experiments and who free-rides in a given network. From this analysis suggestions can be derived on how to structure teams, how to organize information flows in multinational corporations and when a fast diffusion of information is beneficial for society. Moreover, I study monitoring imperfections stemming from intransparencies (or uncertainties) in the patent system and show how these intransparencies affect the incentives of firms to invest in innovative activities. The Federal Trade Commission (2011) emphasizes the importance of a clear patent notice to attain efficiency. However, before implementing policies to increase transparency, it is important to understand the links between intransparencies in

the patent system and R&D investment to be able to assess the consequences of such policies.

Besides monitoring imperfections and different monitoring structures, I analyze how informational externalities interact with different types of payoff externalities - not only within one line of research or sector but also across sectors. The decision of firms in which research lines to invest depends on inter- and intrasectoral R&D spillovers as well as on intellectual property rights and competition. Understanding how informational externalities interact with different levels of inter- and intrasectoral spillovers can tell us how to design intellectual property rights to encourage firms to choose a line of research that is beneficial for society.

To answer these questions I study variants of the exponential bandit model, which was introduced by Keller et al. (2005). This modeling framework is particularly suitable to model fundamental research, where researchers spend considerable time and effort to tackle unsolved problems that might not have a solution and even if a solution exists, it is highly uncertain when it will be discovered. Yet, also other situations feature these types of uncertainty, as e.g., farmers experimenting with a new type of crop or fertilizer where it is unclear whether output can be increased by changing to the new technology. In the first two chapters, I study a discrete time version of the exponential bandit model similar to Heidhues, Rady and Strack (2015). In Chapter 1, I consider a team problem in which one discovery has perfect positive spillovers and study how different network structures affect the incentives of agents to experiment. In Chapter 2, I analyze the impact of inter- and intrasectoral R&D spillovers on the decision of competing firms in which research lines to invest. Finally, in Chapter 3, I investigate a stopping game in continuous time in which heterogeneous (competing) firms invest in R&D and intransparencies in the patent system affect R&D investment.

The first chapter of the thesis analyzes a dynamic model of rational strategic learning in a network. It complements existing literature on learning in networks by providing a detailed picture of the short-run learning dynamics. Agents are located on the nodes of a social network and can observe the actions and outcomes of their direct neighbors immediately, while information generated by agents more far away in the network trav-

els along the links in the network with delay. The complete network, the ring and the star network are compared in terms of their experimentation intensities in equilibrium. The delay in information transmission caused by incomplete network structures induces players to increase own experimentation efforts. As a consequence the complete network can fail to be optimal even if there are no costs for links. This means that in the design of networks there exists a trade-off between the speed of learning and accuracy. Furthermore, the combination of delay and specialization, where only some agents exert effort, can be beneficial for society.

The second chapter analyzes R&D investment with inter- and intrasectoral R&D spillovers, in which the profitability of an invention in one sector increases if there is an innovation in another sector. In the selection of research lines firms face a trade-off; joint research in the same sector increases the probability of an invention, while working in different sectors increases profits in case of two complementary discoveries. Interpreting intrasectoral R&D spillovers as a measure of intellectual property rights I characterize intellectual property rights for which firms select optimally between research lines. In the presence of intersectoral R&D spillovers imitating an invention in the sector benefiting from the spillover is facilitated.

In the third chapter, I analyze how monitoring imperfections in the form of intransparencies in the patent system affect R&D investment. Two firms invest in research of uncertain quality and file an application at the patent office upon an invention. In an intransparent patent system the scope as well as the content of a patent can be unclear. If the patent race is a winner-takes-all competition, the R&D investment of the weaker firm increases in the level of transparency, while the stronger firm's investment only increases if the difference in the firms' R&D productivities is not too severe. In the presence of positive R&D spillovers, R&D investment is non-monotonic in the level of transparency and maximal under full intransparency. An optimal information disclosure policy enables firms to assess whether a new discovery leads to a patent. The risk of litigation or difficulties in identifying relevant patents, however, increase the likelihood that firms invest in R&D. Thus, full transparency is not necessarily optimal provided that the costs of R&D are sufficiently low.

To sum up, in the presence of informational externalities, the incentives of strategic agents to invest in innovative and risky activities are affected by different monitoring structures and monitoring imperfections. A fast and perfect information transmission is not certainly optimal due to the strong incentives of firms to free-ride on the experimentation efforts of others. Incomplete interaction structures or monitoring imperfections such as uncertainties in the patent system can encourage firms to invest in R&D and thereby increase welfare.

The proofs to each chapter can be found in the corresponding sections of the Appendix. All references are collected in the end of the thesis.

Chapter 1

Learning faster or more precisely?

Strategic experimentation in networks

“Some people will never learn anything, for this reason, because they understand everything too soon.” Alexander Pope

1.1 Introduction

The experience of others plays an important role when individuals have to take decisions about alternatives that they cannot perfectly evaluate themselves. This is the case if payoffs associated with one or more alternatives are uncertain. For example, when it comes to the adoption of new technologies it can be ex ante difficult to evaluate whether a new technology will be superior to the status quo. In such situations a person will base his or her decision on own past experiences, ask friends and coworkers about their opinions, and additionally collect information via other sources as for instance, customer reviews on the internet. One way to model learning situations where people have to take decisions under uncertainty is by so called bandit models (see e.g., Bolton & Harris, 1999 or Keller, Rady and Cripps, 2005 [KRC, hereafter]). The idea of these models is that players have to choose between different options (different arms of a bandit machine) under imperfect knowledge of their relative advantage, that is, the outcomes of the arms are uncertain. By playing repeatedly, the agents can learn about the type of the arm, however, this learning or experimentation is costly as future payoffs are discounted. Such bandit

models can provide a framework to discuss different (economic) situations as for example specific problems of product choice, technology adoption, research or innovation.

So far, most models of strategic experimentation assume that agents interact with everyone else in society. That is, each agent can observe or communicate with all other individuals and as actions and payoffs are publicly observable, a common belief about the state of the world prevails. We relax this assumption by letting agents interact directly only with a subset of agents that is determined by the structure of connections in a social network. This extension of the model is thought to better reflect interaction patterns in reality, where without doubt the structure of relationships in shaping beliefs and opinions plays an important role. Empirical work in economics highlights the impact of network structures in labor markets (e.g., regarding information about job vacancies, see Calvo-Armengol and Jackson, 2004) or finds evidence of the importance of interaction patterns in learning about a new technology (see e.g., Conley and Udry, 2010). In general, learning and innovation are influenced by the structures of information exchange between different sources. For example, in the field of research, workshops and conferences bring together researchers from dispersed geographical regions and different fields of specialization to enable exchange of ideas. Similarly, innovation plays an important role for firms to secure competitiveness, and the structure of information exchange between subsidiaries of multinational organizations might be a key to success. For example, Nobel and Birkinshaw (1998) analyze communication patterns between subsidiaries of multinational corporations and find that innovation is associated with comparatively high levels of communication within the firm and outside. Teece (1994) emphasizes the importance of organizational structures that enable an easy flow of communication between business units and guarantee a high speed of learning.

The aim of our model is to provide insight into how the structure of relations influences the evolution of beliefs, decisions and incentives of rational agents who need to acquire information. More precisely, we consider a dynamic model of strategic learning in which individuals can generate own information through experimentation and learn about the experimentation of others. The information generated by other agents diffuses along the links of a social network and only information generated by direct neighbors can be observed immediately. The model is built on a discrete time version of the expo-

nential bandit model by KRC as in Heidhues, Rady and Strack (2015) [HRS hereafter]. Agents can choose between a safe and a risky option. The payoff of the risky option depends on an unknown state of the world which is either good or bad for all players. A good risky option generates high payoffs with positive probability, while a bad risky option never generates a high payoff. Thus, it is not clear whether a high payoff can be obtained and if so when it will occur. These two types of uncertainties are common features of research and innovation. Examples include pharmaceutical firms working on the development of a new drug, mathematicians tackling a Millennium Prize Problem, or farmers experimenting with a new fertilizer as in Conley and Udry (2010). Agents decide between the two options based on their belief, i.e., the probability attached to the good state of the world and their beliefs depend on their observations and hence the interaction structure. This interaction structure will be fixed and imposed on the agents before the game starts.

First, we characterize symmetric equilibria in Markovian strategies in three different network structures, the complete network, the ring and the star network. Further, experimentation intensities in equilibrium are compared across these structures. In a network structure in which agents learn from unobserved players (neighbors of neighbors) with a delay, players increase their experimentation intensity or effort to compensate for the worse possibility to learn from others. Depending on the structure and the belief, agents are able to fully outweigh this loss and thereby keep expected utilities unaltered compared to interaction in a complete network. The agents' strategies depend on their beliefs and there exists an upper cut-off belief above which agents experiment with full intensity and a lower cut-off below which experimentation ceases. These cut-off beliefs depend on the network structure and take into account whether agents still expect information that was generated by unobserved individuals to arrive. Specialization, where only some players experiment, arises in networks where agents are not symmetric with respect to their position, as in the star network. As part of a welfare comparison, we observe a trade-off between interaction structures that enable fast learning and structures in which learning is more precise, or put differently between delay and free-riding. How this trade-off is resolved depends on the discount factor. Additionally to the trade-off between delay and free-riding, we show that also specialization, which occurs naturally

in the star network, motivates agents to exert higher effort. For pessimistic beliefs welfare in the star network is strictly higher than in the ring or complete network because of a combination of delay and specialization.

The chapter contributes to the theory of rational strategic learning in networks and aims to fill the gap between static models and dynamic models that focus on long-run results and conditions for complete learning. Due to the complexity that network settings can create, attention was often restricted to the behavior of myopic or boundedly rational agents to ensure tractability.¹ In a recent contribution Sadler (2014) analyzes a strategic experimentation problem as in Bolton and Harris (1999) in a network setting with boundedly rational agents. In the field of rational learning, there are several examples of Bayesian learning models focusing on asymptotic long-run results and conditions for complete learning or convergence of actions or payoffs (see e.g., Gale and Kariv, 2003, Rosenberg, Solan and Vieille, 2009, Acemoglu and Ozdaglar, 2011, Acemoglu, Bimpikis and Ozdaglar, 2014, Arieli and Mueller-Frank, 2015 and Mossel, Sly and Tamuz, 2015). These results, however, offer little insight into how social relations shape incentives in early stages of the learning process and how this influences expected payoffs.

One closely related paper is Bramoullé and Kranton (2007) [BK hereafter] who investigate a public goods game in a network in a static framework. The authors show that networks can lead to specialization and that this specialization can have welfare benefits, features that are confirmed within our framework. The main difference between the model of BK and our model is that BK consider a static setting. As learning, innovation and research have a dynamic character, a dynamic perspective might be better suited to analyze these processes. Such a perspective yields additional insights concerning the updating rules agents use, the effects of different beliefs within a society that are a consequence of asymmetric positions, and the impact of network structures on the speed and accuracy of learning.

Besides the literature on learning in networks the chapter relates to the growing literature on strategic experimentation. In strategic experimentation models where agents

¹See e.g. Jackson (2008), Chapter 8 or Goyal (2009), Chapter 5 for different types of learning models in a network setting.

can observe the outcomes and actions of others, strong incentives to free-ride on the experimentation effort of others prevent the socially optimal outcome (see e.g., Bolton and Harris, 1999, KRC or HRS). Bonatti and Hörner (2011) show that when actions are private agents do not only experiment too little, but also too late. If actions can be observed but payoffs are private free-riding can be ameliorated when agents have access to cheap talk communication (see HRS). Bimpikis and Drakopoulos (2014) show that (full) efficiency can be obtained in the model of KRC if information is aggregated and released with an optimal delay. An incomplete network structure also causes time lags in the information transmission that can increase experimentation efforts and mitigate free-riding. Additionally to the trade-off between delay and free-riding that is also present in Bimpikis and Drakopoulos (2014) we show that in asymmetric network structures as the star network specialization arises in equilibrium. Furthermore, the combination of delay and specialization can be beneficial for society.

The contribution of the chapter is twofold. First, we introduce delay into a model of rational (and farsighted) learning in networks. Thereby we are able to highlight the importance of the trade-off between delay and free-riding for social learning when the delay is determined by the interaction structure. Second, by adding a network structure to a game of strategic experimentation we show how equilibrium experimentation varies with the interaction structure. As a consequence of asymmetric positions in the network specialization can arise in equilibrium.

The chapter is structured as follows: Section 1.2 introduces the basic model. In Section 1.3 the complete network is analyzed to set up a benchmark case for future comparison. Section 1.4 analyzes a simple incomplete interaction structure, namely a ring, to see how spatial structures change the problem at hand. In Section 1.5, the star network as the simplest irregular network is considered. A welfare analysis is conducted in Section 1.6. Section 1.7 contains a discussion and conclusion. All proofs are relegated to Appendix A.

1.2 Model

First, we describe the underlying bandit model. After that, main concepts of the network structure are outlined and the timing and information structure are specified. With the

help of a short example we briefly show how a network structure affects updating rules. Finally, strategies as well as the equilibrium concept are discussed.

1.2.1 A two-armed bandit model

The model is based on the two-armed exponential bandit model as described by KRC or more specifically the discrete time version thereof by HRS. There are agents $i \in N$ and we denote the cardinality of N by n . Time is discrete and players discount future payoffs by a common discount factor $\delta \in (0, 1)$.

Players can experiment with an uncertain technology. To be more precise, in each time period $t = 1, \dots, T$ each agent is endowed with one unit of a perfectly divisible resource (e.g., effort or money) that can be allocated between a risky and safe technology (which correspond to the risky and the safe arm of two-armed bandit machine). Let $\phi_{i,t} \in [0, 1]$ denote the fraction of the unit resource that is allocated to the risky arm and $1 - \phi_{i,t}$ is allocated to the safe arm. Subsequently we will refer to $\phi_{i,t}$ as experimentation effort or experimentation intensity, however, $\phi_{i,t}$ can be also interpreted as R&D investment.

The safe arm yields a fixed deterministic payoff normalized to 0. The risky arm (denoted by R) yields an uncertain reward $X_{i,t} \in \{0, X_H\}$ with $X_H > 0$. The distribution of the risky payoffs is independent across players and time and only depends on the state of the world, which is either good ($\theta = 1$) or bad ($\theta = 0$) for all players. In the bad state of the world the probability of obtaining a high payoff X_H is zero. If the state of the world is good, R yields a high payoff X_H with positive probability. This probability is proportional to the level of effort invested and given by $\phi_{i,t}\pi$, where $\pi > 0$. Consequently, the first high payoff realization (also called a breakthrough) perfectly reveals that the risky arm is good. Further, investing into the risky arm is costly and costs are proportional to effort invested. That is, an agent choosing effort level $\phi_{i,t}$ pays costs of $\phi_{i,t}c$ at time t where $c > 0$. Thus, $\phi_{i,t}$ is a pure action that on the one hand determines the costs an agent has to pay for experimenting with the risky technology and on the other hand scales the probability of success.

In any given time period the expected risky payoff conditional on the state of the world equals $-\phi_{i,t}c$ if the state is bad and $\phi_{i,t}\pi X_H - \phi_{i,t}c$ if the state is good. In what

follows we will denote the expected low payoff $-c$ by E_0 and the expected high payoff $\pi X_H - c$ by E_1 . Additionally to the fact that $E_0 < 0$ we assume that $E_1 > 0$, which means that it is optimal for all players to use the risky arm if $\theta = 1$ and use the safe arm if $\theta = 0$.

Players hold a belief p about the risky arm being good, and it is assumed that they start with a common prior. Agents influence each other only through the impact of their action on the belief of others, meaning there are only informational externalities and no payoff externalities. In the model of HRS and KRC all players interact with everyone else in the society and hence agents hold a common posterior belief. This will no longer be true when players interact only with a subset of society. A player's belief depends on whether she learns about a breakthrough. Once she does, her uncertainty about the type of the arm is resolved and the posterior belief jumps to 1. As long as agents experiment without learning about a breakthrough, beliefs are updated according to Bayes' rule and decrease (no news is bad news). Players are said to experiment if they use the risky arm before knowing its type.

1.2.2 Introducing a network structure

Given a set of nodes N (representing individuals), a network or graph g is an $n \times n$ interaction matrix that represents the relationships in the society. The typical element of g is denoted by $g_{ij} \in \{0, 1\}$. If $g_{ij} = 1$, a link between i and j exists and implies that agent i observes agent j 's actions and outcomes immediately without delay. The matrix is symmetric ($g_{ij} = g_{ji}$), meaning links are undirected, and always has 1 on the main diagonal (every individual can observe her own actions and outcomes, i.e., $g_{ii} = 1$ for all i). The structure of relations is assumed to be common knowledge. If a link between two individuals exists, those agents are considered to be neighbors. The *neighborhood* of agent i is denoted by N_i and defined as $N_i(g) = \{j \neq i : g_{ij} = 1\}$.

Subsequently a fixed interaction structure g will be imposed. The game is analyzed in three different network structures: the complete network² as a benchmark case; the ring, an incomplete but regular³ structure; and the star network with one player in the center and all other $n - 1$ players only connected to the central player (see Figure 1.1).

²A complete network is a network in which every agent is connected to everyone else.

³Regular networks are networks where all players have the same number of neighbors.




Figure 1.1: Network structures for $n = 6$.

1.2.3 Timing and information structure

The timing of the game is as follows: agents start in $t = 1$ with a common prior belief p_1 . Each agent chooses an experimentation intensity or effort $\phi_{i,1} \in [0, 1]$, determining how much effort is invested in the risky option. At the end of $t = 1$ players observe their own outcomes as well as the actions and outcomes of their neighbors and update their prior accordingly to $p_{i,2}$. Those agents who have not observed a success choose $\phi_{i,2} \in [0, 1]$. That is, $\phi_{i,2}$ is the experimentation effort conditional on not having observed a breakthrough in $t = 1$. Agents then observe outcomes and actions in their neighborhood and exchange verifiable reports about previous experiments by unobserved agents, i.e., in $t = 2$ agent i knows $\phi_{m,1}$ as well as $X_{m,1}$ for all agents $m \in N_j \setminus N_i$ where $j \in N_i$. This process of information transmission continues in the subsequent periods. Agents do not need to draw inferences from the actions of their neighbors as they receive this information through the report. The exchange of reports in our game takes place automatically and agents can neither choose whether they want to exchange reports nor the information these reports contain.⁴ That is, agents can only choose their experimentation effort $\phi_{i,t}$.

Formally, agent i 's information at a given point in time t consists of

$$\mathcal{I}_{i,t} = \{\mathbf{H}_{i,t}, \mathbf{r}_{i,t}\},$$

where $\mathbf{H}_{i,t} = \{\phi_{i,1}, X_{i,1}, \dots, \phi_{i,t}, X_{i,t}\}$ is the complete history of actions and outcomes for agent i up to time t and $\mathbf{r}_{i,t} = (r_{i,1}, \dots, r_{i,t})$ is the history of reports agent i received. Each

⁴If agents are allowed to freely choose any message, they may find it optimal to report a breakthrough although there was none in order to induce additional experiments. See HRS for a strategic experimentation game in which payoffs are privately observed and agents exchange cheap talk messages.

element $r_{i,t}$ is a vector that contains for each agent $j \in N_i$ the history $\mathbf{H}_{j,t}$ up to this point in time as well as the reports j received up to $t - 1$, i.e., $\mathbf{r}_{j,t-1}$.

For $t > T$ agents cannot experiment anymore, i.e., they are restricted to the safe option ($\phi_{i,t} = 0$ for all $t > T$) if they did not learn about a success up to time T . Information diffusion still takes places after T and agents can switch to the risky option in case they learn that the state of the world is good.⁵ After learning about a breakthrough the agent continues to use the risky option forever.

As soon as the network is incomplete at least some of the agents do not possess complete information about (past) actions and payoffs of others. Consequently, when interacting with their neighbors, agents obtain information through them about (past) actions and payoffs of unobserved agents and use this information to make inferences about the true state of the world. In incomplete networks the probability of learning about a breakthrough at a given point in time depends on the entire structure of relations, and information about a breakthrough will travel along the paths in the network. This implies that players will not necessarily hold a common belief about the state of the world. We will illustrate the impact of the network structure on the updating of beliefs with the help of a short example.




Figure 1.2: The star network, $n = 3$.

Example 1 There are three agents $i = 1, 2, 3$, whose connections can be described by the following interaction matrix (see Figure 1.2)

$$g = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}.$$

⁵Restricting the experimentation period allows us to solve for the equilibrium experimentation effort by backward induction. Besides this, in many situations arbitrarily long experimentation might not be possible. For example, research funds are often granted for a certain time period and not prolonged in the absence of a success or new (risky) technologies with more promising success probabilities will cause agents to stop experimenting with a less attractive alternative.

As $g_{2j} = 1$ for all $j \in N$, agent 2 has complete information and can observe all actions and payoffs at any point in time. The other two agents only observe agent 2 and their own actions and payoffs, and they receive information through agent 2. A success by agent 2 immediately reveals to everyone that the risky arm is good. If agent 1 has a breakthrough, only 1 and 2 know about it. However, agent 3 learns about the breakthrough through agent 2 so that agent 3 knows about it one period later. As long as there is no breakthrough, agents update their beliefs depending on how many unsuccessful experiments they learn about. That is, player 2 updates her belief according to

$$p_{2,t+1} = \frac{p_{2,t} \prod_{i=1}^3 (1 - \phi_{i,t}\pi)}{p_{2,t} \prod_{i=1}^3 (1 - \phi_{i,t}\pi) + 1 - p_{2,t}}$$

if no breakthrough occurs, where $\prod_{i=1}^3 (1 - \phi_{i,t}\pi)$ reflects the experiments conducted by 2 and her neighbors. The numerator is the probability of not observing a breakthrough on a good risky arm and the denominator gives the total probability of not observing a breakthrough. In case of a breakthrough the posterior jumps to 1. Player 1 updates her belief according to

$$p_{1,t+1} = \frac{p_{1,t} \prod_{i=1}^2 (1 - \phi_{i,t}\pi)(1 - \phi_{3,t-1}\pi)}{p_{1,t} \prod_{i=1}^2 (1 - \phi_{i,t}\pi)(1 - \phi_{3,t-1}\pi) + 1 - p_{1,t}},$$

if there is no breakthrough. At time t agent 1 observes the outcome of her own experiment as well as agent 2's experiment. While agent 1 does not observe agent 3's current experiment she gets informed about the experiment performed in $t-1$, which is captured by the term $(1 - \phi_{3,t-1}\pi)$. Agent 3's belief at time t is derived analogously.

1.2.4 Strategies and equilibrium concept

Players are fully rational and maximize their expected payoffs. In period t , agent i obtains a payoff of $\phi_{i,t}X_{i,t}$, with $X_{i,t} \in \{0, X_H\}$, and player i 's total expected (normalized and discounted) payoff is given by

$$(1 - \delta)E \left[\sum_{t=1}^{\infty} \delta^{t-1} \phi_{i,t} X_{i,t} \right], \quad (1.1)$$

where the expectation is taken w.r.t. $p_{i,t}$ and $\phi_{i,t}$. Players are restricted to pure strategies that only depend on payoff relevant information. In the complete network this corresponds to the Markov perfect equilibrium with the common posterior belief as the state variable. In incomplete network structures the agents' strategies depend on the state of the world as described by the past beliefs of the agents and additional information about the network structure and the positions of the agents in the network. In what follows we restrict attention to equilibria in which agents who are symmetric with respect to their position in a network use symmetric strategies.

1.3 The empty and the complete network

Before we analyze experimentation in incomplete networks we first explain how the experimentation problem is solved by a single agent. After that, we look at the model with n agents, where each individual can observe everyone else. Expressed in terms of networks this corresponds to the empty and the complete network.

1.3.1 The empty network

The agent maximizes expected payoffs by choosing a sequence of actions $\{\phi_{i,t}\}_{t=1}^T$ in $t = 1$, where $\phi_{i,t+1}$ is the effort chosen in $t+1$ conditional on not having observed a breakthrough until time t . The expected payoff at time t can be written recursively as

$$U(p_t) = \phi_{i,t}(1 - \delta)E_{p_t} + \delta E_1 p_t \phi_{i,t} \pi + \delta(1 - p_t \phi_{i,t} \pi)U(p_{t+1}), \quad (1.2)$$

with $E_p = E_1 p + (1 - p)E_0$. The agent's posterior belief is given by

$$p_{i,t+1} = \frac{p_{i,t}(1 - \phi_{i,t}\pi)}{1 - p_{i,t}\phi_{i,t}\pi} \quad (1.3)$$

if she does not observe a breakthrough. The posterior jumps to 1 after a success. In expression (1.2) the first part, $\phi_{i,t}(1 - \delta)E_{p_t}$, is the expected and normalized current payoff the agent obtains in period t by exerting effort $\phi_{i,t}$. A good risky arm generates a payoff of E_1 , while a bad risky arm gives E_0 . The remaining terms represent the discounted expected continuation payoff. The continuation payoff is E_1 with the probability $p_t \phi_{i,t} \pi$ that the risky arm is good and a breakthrough occurs. If the agent does not observe a

success she can experiment again in $t+1$. The probability of not observing a breakthrough consists of the probability that the risky arm is bad, $1 - p_t$, and the probability that it is good, but the agent nevertheless did not have a breakthrough, $p_t(1 - \phi_{i,t}\pi)$.

For a given prior the expected (normalized) payoff can also be written as

$$U(p_1) = (1 - \delta) \left\{ (1 - p_1)E_0 \sum_{t=1}^T \delta^{t-1} \phi_{i,t} + p_1 E_1 \left(1 + \frac{\delta\pi}{1 - \delta} \right) \sum_{t=1}^T \delta^{t-1} \phi_{i,t} \prod_{r=1}^{t-1} (1 - \phi_{i,r}\pi) \right\}, \quad (1.4)$$

and consists of two parts. First, the player obtains a low expected payoff E_0 in every period she experiments, if the state of the world is bad. Second, if the state of the world is good, the agent may obtain a high payoff. More precisely, $\prod_{r=1}^{t-1} (1 - \phi_{i,r}\pi)$ is the probability of not having had a success up to period t and $E_1 \frac{\delta\pi}{1 - \delta}$ is the discounted expected continuation payoff after a breakthrough. The expected payoff is linear in the prior p_1 and in effort, i.e., linear in each element of the sequence $\{\phi_{i,t}\}_{t=1}^T$. Linearity in effort implies that in case expected payoffs of experimenting in period t are strictly positive, the risky option is strictly preferred to the safe option and the agent allocates the entire resource to the risky option, i.e., $\phi_{i,t} = 1$. If expected payoffs of experimenting are negative it is optimal not to experiment and the agent chooses $\phi_{i,t} = 0$. In case expected payoffs of experimenting are equal to zero, the agent is indifferent between all $\phi_{i,t} \in [0, 1]$.

Expected payoffs are increasing in the belief and the belief decreases over time in the absence of a success. Thus, the agent experiments for optimistic beliefs and stops for pessimistic beliefs. At time T the continuation payoff $U(p_{T+1}) = 0$ if there was no success so far and the payoff of experimenting one more time, $(1 - \delta)E_{p_T} + \delta\pi E_1 p_T$, is higher than the expected payoff of stopping for any belief p_T greater or equal to

$$p^a = \frac{(1 - \delta)|E_0|}{(1 - \delta)[|E_0| + E_1] + \delta E_1 \pi}, \quad (1.5)$$

where a stands for autarky. Using backward induction we can solve for ϕ_{T-1} given ϕ_T and show that it is never optimal to postpone experimentation as future payoffs are discounted. This means that for pessimistic prior beliefs the agent never experiments, for more optimistic priors the agent experiments only in $t = 1$, and so forth, where experimenting one more time at time t is optimal if $p_t \geq p^a$ (see Figure 1.3).

Lemma 1.1. *Each period a single agent either uses the safe arm exclusively or the risky arm*

exclusively, and for $t \leq T$ her action only depends on the belief in the given period, that is,

$$\phi_{i,t}^a = \begin{cases} 1 & \text{for } p_t \geq p^a, \\ 0 & \text{otherwise.} \end{cases}$$




Figure 1.3: Expected payoffs of experimenting for the single agent. $U(1,0)$ represents the expected payoff of experimenting in $t = 1$ only and $U(1,1)$ the expected payoff of experimenting in $t = 1, 2$ only. The point of intersection of $U(1,0)$ with $U(1,1)$, denoted by $p_1 = p^*$, is the prior belief at which $p_2 = \frac{p^*(1-\pi)}{1-p^*\pi} = p^a$.

The single agent's strategy depends only on the current belief, as this belief captures all payoff-relevant information. Agents stop experimenting at $p^a > 0$, which means that it is possible that they abandon the risky arm although it is good. The cut-off belief p^a decreases in δ , which means that as agents are getting more patient, complete learning becomes more likely. That is, the final posterior belief in case all experiments fail is smaller and hence the probability of mistakenly switching from the risky to the safe arm although the risky arm is good decreases. We will subsequently refer to this final posterior also as the *precision* of learning and say that learning is more precise the lower this final posterior.

1.3.2 The complete network

In a complete network each player maximizes her expected utility given her belief and the strategies of the other players. Let $\tilde{\phi}_{-i,t}$ denote $\prod_{j=1}^n (1 - \phi_{j,t}\pi)$ for all $j \neq i$. That is, $\tilde{\phi}_{-i,t}$ represents the probability that none of the agents $j \neq i$ has a breakthrough at time t .

Agent i 's expected payoff at time t is given by

$$U_i(p_t) = \phi_{i,t}(1 - \delta)E_{p_t} + \delta E_1 p_t (1 - \tilde{\phi}_{-i,t}(1 - \phi_{i,t}\pi)) + \delta \left(1 - p_t(1 - \tilde{\phi}_{-i,t}(1 - \phi_{i,t}\pi))\right) U_i(p_{t+1}), \quad (1.6)$$

where

$$p_{t+1} = \frac{p_t \tilde{\phi}_{-i,t}(1 - \phi_{i,t}\pi)}{p_t \tilde{\phi}_{-i,t}(1 - \phi_{i,t}\pi) + 1 - p_t}. \quad (1.7)$$

The continuation payoff of agent i now also depends on the actions of the other players. Otherwise, the problem is similar to the single agent, that is, payoffs are linear in the belief and in effort. Proposition 1.1 describes the optimal experimentation effort in a symmetric equilibrium.

Proposition 1.1. *In the symmetric equilibrium in a complete network with n agents*

(i) *the common strategy for $t \leq T$ is given by*

$$\phi_t^c = \begin{cases} 1 & \text{for } p_t \in [\bar{p}^c, 1], \\ \frac{1}{\pi} - \frac{1}{\pi} \left(\frac{(1-\delta)|E_0|}{\delta E_1 \pi p_t} - \frac{(1-\delta)(|E_0|+E_1)}{\delta E_1 \pi} \right)^{\frac{1}{n-1}} & \text{for } p_t \in (p^a, \bar{p}^c), \\ 0 & \text{for } p_t \in [0, p^a], \end{cases} \quad (1.8)$$

where

$$\bar{p}^c = \frac{(1 - \delta)|E_0|}{(1 - \delta)[|E_0| + E_1] + \delta E_1 \pi (1 - \pi)^{n-1}};$$

(ii) *there is at most one time period in which $\phi_t^c \in (0, 1)$.*

The players' strategies depend only on the common belief p_t as this belief captures all payoff relevant information (actions and payoff realizations) up to time t . An agent's experimentation effort at time t depends on the current (and future) actions of the other players which are determined by the current belief. Any deviation will be captured by the belief and we assume that after a deviation every agent plays optimally given the common posterior.

There exists an interval of beliefs such that in a symmetric equilibrium players simultaneously use both arms. In this interval ϕ_t^c is chosen such that agents are indifferent between the risky and the safe arm, or to be more precise between all $\phi_t \in [0, 1]$. There exists an upper cut-off belief, \bar{p}^c , which is the belief above which agent i experiments with intensity 1 even if all others also experiment with full intensity. Starting from \bar{p}^c agents

decrease their experimentation intensity as the belief decreases, down to the point where $\phi_t^c = 0$, which holds for any belief below p^a . Figure 1.4 depicts the equilibrium strategy.




Figure 1.4: Experimentation effort in a symmetric equilibrium in the complete network ($\pi = 0.2$, $\delta = 0.9$, $n = 12$, $E_1 = 1$, $E_0 = -1$, $p^a \approx 0.26$ and $\bar{p}^c \approx 0.46$).

Several features of the equilibrium experimentation strategy are worth noting. First, there is at most one time period in which agents simultaneously use both arms. In fact, for any $n \geq 2$, n failed experiments from \bar{p}^c generate a posterior belief below p^a , and the common effort $\phi^c(p_t)$ at beliefs $p_t \in (p^a, \bar{p}^c)$ also causes the posterior to fall below the single agent cut-off if there is no success. Second, agents do not have an incentive to delay any experiments: if $\phi_t^c > 0$, then $\phi_{t-1}^c = 1$. Third, the upper cut-off \bar{p}^c is increasing in n , whereas the lower threshold is given by p^a . Social optimality requires experimentation beyond the single agent cut-off, p^a , since agents benefit from each others experimentation effort. However, agents do not experiment below p^a and even stop experimenting with full intensity earlier as the number of agents increases, which is a particularly stark manifestation of the free-riding effect (see also HRS or KRC).

1.4 The ring network

Let us now turn to the strategic experimentation problem when agents are located on a ring. In the ring network every agent has two direct neighbors. As players are symmetric

there again exists a symmetric equilibrium. The underlying structure is illustrated in Figure 1.1b for $n = 6$.

Expected payoffs in the ring network for a given prior and strategy profile are

$$U_i(p_1) = \phi_{i,1}(1 - \delta)E_{p_1} + \delta E_1 p_1 [1 - (1 - \phi_{j,1}\pi)^2(1 - \phi_{i,1}\pi)] \\ + \delta [1 - p_1 + p_1(1 - \phi_{j,1}\pi)^2(1 - \phi_{i,1}\pi)] U_i(p_2);$$

as we are solving for symmetric equilibria, we are assuming here that all agents $j \neq i$ use the same strategy. The term

$$U_i(p_2) = \phi_{i,2}(1 - \delta)E_{p_2} + \delta E_1 p_2 [1 - (1 - \phi_{j,2}\pi)^2(1 - \phi_{i,2}\pi)(1 - \phi_{j,1}\pi)^2] \\ + \delta [1 - p_2 + p_2(1 - \phi_{j,2}\pi)^2(1 - \phi_{i,2}\pi)(1 - \phi_{j,1}\pi)^2] U_i(p_3),$$

as well as $U_i(p_3)$, $U_i(p_4)$ and so on, is determined by the information about past experiments traveling through the network. In $t = 1$ the agents start with a prior belief p_1 , choose their experimentation intensity $\phi_{i,1}$ and receive their payoffs. Then each agent either knows that the state of the world is good if there was a breakthrough in her neighborhood, or she chooses her optimal experimentation intensity $\phi_{i,2}$ based on her updated belief $p_{i,2}$. That is, with the probability that at least one experiment agent i learns about in $t = 1$ is successful, $p_1[1 - (1 - \phi_{j,1}\pi)^2(1 - \phi_{i,1}\pi)]$, she gets a continuation payoff of E_1 from the next period onwards. These are the two experiments of the neighbors as well as the own experiment. In case all these experiments were unsuccessful, she and her neighbors can experiment again in $t = 2$. Further there is the chance that neighbors of neighbors had a breakthrough in $t = 1$ about which agent i will learn in $t = 2$. That is, in $t = 2$ the agent receives information about the outcome of the first period experiment of the neighbors of neighbors. In $U_i(p_2)$, the factor $(1 - \phi_{j,1}\pi)^2$ in $1 - (1 - \phi_{j,2}\pi)^2(1 - \phi_{i,2}\pi)(1 - \phi_{j,1}\pi)^2$ represents the experiments of neighbors of neighbors in $t = 1$, $(1 - \phi_{j,2}\pi)^2$ the two experiments of the direct neighbors in $t = 2$, and $1 - \phi_{i,2}\pi$ the own experiment in $t = 2$. This process continues until either all agents stopped and all information has reached agent i , or every agent knows that the state is good. The information transmission takes the longer the more players there are. The posterior belief of an agent, who did not observe a

breakthrough so far, for symmetric actions is given by

$$p_{i,t+1} = \frac{p_{i,t}(1 - \phi_{i,t}^r \pi)^3 \prod_{l=1}^{\min\{t-1, d-1\}} (1 - \phi_{i,t-l}^r \pi)^2}{1 - p_{i,t} + p_{i,t}(1 - \phi_{i,t}^r \pi)^3 \prod_{l=1}^{\min\{t-1, d-1\}} (1 - \phi_{i,t-l}^r \pi)^2}, \quad (1.9)$$

where d denotes the diameter of the network.

In contrast to the complete network, equilibrium cut-off beliefs vary with time. This can be ascribed to the fact that after one period of experimentation, information is traveling through the network and agents anticipate that this information will reach them. For example, in $t = 2$ the lower cutoff belief \underline{p}_2^r is above p^a , because of the two experiments conducted by neighbors of neighbors in $t = 1$ that agent i learns about in $t = 2$. In order to analyze the equilibrium behavior of the agents we introduce the expression I_t^r . I_t^r represents the difference in expected payoffs from experimenting with full intensity and not experimenting at all for symmetric actions of the other players in period t with no experimentation in $t + 1$. This means that $I_t^r > 0$ implies that payoffs from experimenting are higher than payoffs from not experimenting in t and at $I_t^r = 0$ agents are indifferent. Proposition 1.2 describes the symmetric equilibrium in the ring network. The expression for I_t^r and the equilibrium cut-off beliefs can be found in Appendix A.

Proposition 1.2. *In the symmetric equilibrium in the ring network for $t \leq T$ each player chooses the following action:*

- $\phi_t^r = 1$ for $p_{i,t} \in [\bar{p}_t^r, 1]$,
- $\phi_t^r = 0$ for $p_{i,t} \in [0, \underline{p}_t^r]$, where $\underline{p}_1^r = p^a$ and,
- $\phi_t^r \in (0, 1)$ is defined uniquely by the root of I_t^r on $[0, 1]$ for $p_{i,t} \in (\underline{p}_t^r, \bar{p}_t^r)$ and chosen such that each player is indifferent between all $\phi_t \in [0, 1]$ for symmetric actions of the other players. There is at most one time period in which $\phi_t^r \in (0, 1)$.

The agents' actions depend on their belief about the state of the world. In contrast to the complete network, agents in the ring network do not hold a common posterior belief. The threshold beliefs of the agents depend on the state of the world in period $t - d + 1$, where the state of the world is described by a vector of beliefs and corresponding network

positions for each agent. The diameter of the network, d , determines the maximum delay in the network. Consequently at time t the belief (together with the network position) of each agent in the network at time $t - d + 1$ is common knowledge and all information up to time $t - d + 1$ can be ignored.⁶ The best response of agent i depends on the state of the world at time $t - d + 1$ as well as on her private information, i.e., her own private belief at time t , the beliefs of her direct neighbors in $t - 1$ and so forth. When anticipating the actions of the other players, each player assumes that everyone plays a best response given the commonly known state in $t - d + 1$.⁷

As can be seen in Proposition 1.2 the lower cut-off below which experimentation ceases in the first period is equal to the single agent cut-off. The upper cut-off in the ring

$$\bar{p}_1^r = \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + \delta E_1 \pi [(1 - \pi)^2 - [1 - (1 - \pi)^2] \sum_{t=1}^{d-1} \delta^t (1 - \pi)^{2t}]}$$

is smaller than the upper cutoff in the complete network \bar{p}^c with the difference $\bar{p}^c - \bar{p}_1^r$ monotonically increasing in n . This difference increases in the number of players because information needs longer to be transmitted in the ring. The longer agents have to wait for information, the more likely they will find it optimal to experiment themselves in the meantime.

We are interested in the difference between the complete network and the ring in terms of experimentation effort in equilibrium. Proposition 1.3 below shows that in the ring network effort is never lower than in the complete network. In $t = 1$ for high beliefs all agents in both networks experiment, for pessimistic beliefs no one experiments and for intermediate beliefs where agents use both options, the experimentation intensity is higher in the ring. This shows that agents compensate a worse possibility to learn from others through increased own effort. Figure 1.5 illustrates this finding.

If all experiments in $t = 1$ fail, beliefs in the two networks in $t \geq 2$ are different as agents in the complete network are already more pessimistic. Taking the difference in

⁶For $t < d$ it is the state of the world in $t = 1$ that is commonly known by all agents.

⁷Note that this implies that the threshold belief does in fact not depend on time and is constant for any $d \leq t \leq T$.




Figure 1.5: Equilibrium experimentation effort in a ring network (dashed line) and the complete network (solid line) in $t = 1$ ($\pi = 0.2$, $\delta = 0.9$, $n = 12$, $E_1 = 1$, $E_0 = -1$, $p^a \approx 0.26$, $\bar{p}^c \approx 0.46$ and $\bar{p}_1^r \approx 0.34$).

posterior beliefs into account, it can be shown that experimentation effort in the ring and the complete network is either the same, or that effort is higher in the ring. To compare efforts in $t \geq 2$ across different networks, we express beliefs in terms of p_t^c . This means that we make use of the fact that in equilibrium the relationship between the posterior beliefs in the two networks is given by

$$p_t^c = \frac{p_t^r (1 - \pi)^{y_{t-1}}}{p_t^r (1 - \pi)^{y_{t-1}} + 1 - p_t^r}, \quad (1.10)$$

where

$$y_t = (n - 3)t - 2 \sum_{x=1}^{\min\{t-1, d-1\}} (t - x) \quad (1.11)$$

is the difference between the number of experiments an agent in the complete network observed and the number of experiments an agent in the ring observed until time t (if everyone experiments until time t with full intensity).⁸

⁸Note that the expression for y_t in (1.11) describes the situation when n is odd. For an even number of players the term $1_{\{t > d\}}(t - d)$ has to be added.

Proposition 1.3. *Experimentation intensities in the symmetric equilibrium in the ring network are at least as high as in the symmetric equilibrium in the complete network. More precisely,*

$$\begin{aligned}\phi_t^r &> \phi_t^c \quad \text{for } p_t^c \in (\underline{p}_t^r, \bar{p}^c), \\ \phi_t^r &= \phi_t^c \quad \text{for } p_t^c \in [0, \underline{p}_t^r] \cup [\bar{p}^c, 1],\end{aligned}$$

for $t \leq T$, where

$$\underline{p}_t^r = \frac{\underline{p}_t^r (1 - \pi)^{y_{t-1}}}{\underline{p}_t^r (1 - \pi)^{y_{t-1}} + 1 - \underline{p}_t^r} \leq p^a.$$

Over certain intervals of beliefs agents in the ring network exert higher effort than agents in the complete network. Further, agents in the ring network experiment in $t \geq 2$ at beliefs for which the posterior after tn failed experiments is below p^a . That is, \underline{p}_t^r is the posterior belief in the complete network p_t^c at time t that corresponds to a posterior of \underline{p}_t^r in the ring network. If information arrives with delay, agents might be better off experimenting themselves instead of waiting for information generated by others. However, as this information will eventually reach them, the final posterior belief in the ring network can be more pessimistic than in the complete network. That is, the probability of mistakenly abandoning a good risky project decreases and learning is more accurate. This is in line with the finding of Bimpikis and Drakopoulos (2014) that delaying information revelation increases experimentation. The *speed of learning*, measured by the number of time periods until information has traveled to every node in the network, decreases due to the incomplete network structure. Free-riding, however, is reduced as players increase effort over certain intervals of beliefs when information arrives with a delay.

1.5 The star network

To obtain a better understanding of the role of different interaction structures, we now turn to the star network to explore the impact of asymmetric positions on equilibrium experimentation. In the star network one player, called the hub, is located in the center and has a link to each of the other $n - 1$ players. The players at exterior positions, also called peripheral players, are only connected to the hub. The expected payoff of the hub

is given by

$$U^h(p_t^h, p_t^s) = \phi_t^h(1 - \delta)E_{p_t^h} + \delta E_1 p_t^h \left(1 - (1 - \phi_t^h \pi) \prod_{i \neq h} (1 - \phi_{i,t}^s \pi) \right) + \delta \left(1 - p_t^h + p_t^h (1 - \phi_t^h \pi) \prod_{i \neq h} (1 - \phi_{i,t}^s \pi) \right) U^h(p_{t+1}^h, p_{t+1}^s),$$

where

$$p_{t+1}^h = \frac{p_t^h (1 - \phi_t^h \pi) \prod_{i \neq h} (1 - \phi_{i,t}^s \pi)}{1 - p_t^h + p_t^h (1 - \phi_t^h \pi) \prod_{i \neq h} (1 - \phi_{i,t}^s \pi)},$$

and

$$p_{t+1}^s = \frac{p_t^s (1 - \phi_t^h \pi) (1 - \phi_{i,t}^s \pi) \prod_{j \neq i, h} (1 - \phi_{j,t-1}^s \pi)}{1 - p_t^s + p_t^s (1 - \phi_t^h \pi) (1 - \phi_{i,t}^s \pi) \prod_{j \neq i, h} (1 - \phi_{j,t-1}^s \pi)}.$$

For the agents in the periphery we have

$$U_i^s(p_t^s, p_t^h) = \phi_{i,t}^s (1 - \delta) E_{p_t^s} + \delta E_1 p_t^s \left(1 - (1 - \phi_t^h \pi) (1 - \phi_{i,t}^s \pi) \prod_{j \neq i, h} (1 - \phi_{j,t-1}^s \pi) \right) + \delta \left(1 - p_t^s + p_t^s (1 - \phi_t^h \pi) (1 - \phi_{i,t}^s \pi) \prod_{j \neq i, h} (1 - \phi_{j,t-1}^s \pi) \right) U_i^s(p_{t+1}^s, p_{t+1}^h).$$

Players are no longer symmetric and hence an equilibrium in which all players use the same strategy does not exist. In Proposition 1.4 we construct an equilibrium where peripheral players use symmetric strategies and the hub exerts less effort than agents in a symmetric equilibrium in the complete network. More precisely, the hub exerts full effort until \bar{p}^c and does not experiment at all for beliefs below \bar{p}^c . The peripheral players exert higher effort than the players in the complete network in equilibrium.

In the star network the cutoff beliefs in $t = 1$ differ from later periods in which the peripheral agents anticipate the arrival of information generated by unobserved players in the previous period. For $t \geq 2$ in every period the peripheral agents learn about $n - 2$ experiments with one period delay. The diameter of the star network $d = 2$ and consequently the state of the world in $t - 1$ is commonly known by all players. Proposition 1.4 describes the equilibrium where I_t^s is the respective counterpart to I_t^r for the peripheral players in the star network. The expressions for the cut-off beliefs and I_t^s can again be found in Appendix A.

Proposition 1.4. *The strategic experimentation game in the star network where peripheral agents use symmetric strategies has an equilibrium in which the experimentation intensity of the hub for $t \leq T$ satisfies*

$$\phi_t^h = \begin{cases} 1 & \text{for } p_t \in [\bar{p}^c, 1], \\ 0 & \text{otherwise.} \end{cases}$$

For the peripheral players equilibrium experimentation intensities in $t = 1$ are

- $\phi_1^s = 1$ for $p_1 \in [\bar{p}_1^s, 1]$,
- $\phi_1^s = 0$ for $p_1 \in [0, p^a]$,
- $\phi_1^s \in (0, 1)$ is defined uniquely for $p_1 \in (p^a, \bar{p}_1^s)$ by the root of I_1^s on $[0, 1]$ and is chosen such that each peripheral player is indifferent between all $\phi_1 \in [0, 1]$ for symmetric actions of the other peripheral players and $\phi_1^h = 0$.

Experimentation intensities for $t = 2, \dots, T$ are

- $\phi_t^s = 1$ for $p_t \in [\bar{p}_2^s, 1]$,
- $\phi_t^s = 0$ for $p_t \in [0, \underline{p}_2^s]$,
- $\phi_t^s \in (0, 1)$ is defined uniquely for $p_t \in (\underline{p}_2^s, \bar{p}_2^s)$ by the root of I_2^s on $[0, 1]$ and is chosen such that each peripheral player is indifferent between all $\phi_t \in [0, 1]$ for symmetric actions of the other peripheral players and $\phi_t^h = 0$.

In equilibrium there is at most one time period in which $\phi_t^s \in (0, 1)$.

Agents are no longer in symmetric positions and the hub faces a different problem than the peripheral players. In particular, the central player is completely informed about all experiments like in a complete network. Hence, it is optimal for the hub to experiment with full intensity for any belief above \bar{p}^c . If the peripheral players exert higher effort than agents in the complete network, that is if $\phi_t^s > \phi_t^c$, the best response for the hub is not to experiment at all. As it can be shown that $\bar{p}_1^s < \bar{p}^c$, we know that in the interval $[\bar{p}_1^s, \bar{p}^c)$ the hub does not experiment. Further, the best response for the peripheral players to $\phi_1^h = 0$ in (p^a, \bar{p}^c) is to exert higher effort than agents in the complete network. We will refer to a strategy profile in which some agents experiment while others do not exert any effort

also as *specialization*.⁹ For priors above or below the interval (p^a, \bar{p}^c) there will be full or no experimentation respectively. This is illustrated in Figure 1.6.




Figure 1.6: Equilibrium experimentation effort of the peripheral players in the star network (bold dotted line), the central player in the star network (dashed line) and in a complete network (solid line) in $t = 1$ ($\pi = 0.2$, $\delta = 0.9$, $n = 12$, $E_1 = 1$, $E_0 = -1$, $p^a \approx 0.26$, $\bar{p}^c \approx 0.46$ and $\bar{p}_1^s \approx 0.42$).

Remark 1. For some values of the model parameters there exists a second equilibrium in which the peripheral players use symmetric strategies. In $t = 1$ in this equilibrium the hub exerts full effort for beliefs $p_1 \in [p^a, \bar{p}_1^s) \cup [\bar{p}^c, 1]$ and no effort for $p_1 \in [0, p^a) \cup [\bar{p}_1^s, \bar{p}^c)$. This means that the effort of the hub is non-monotonic in the belief. The peripheral agents exert full effort for beliefs above \bar{p}_1^s . For any belief in $[p^a, \bar{p}_1^s)$ their experimentation intensity is lower than the experimentation intensity that makes the hub indifferent, that is, $\phi_1^s < \phi_1^c$. The second equilibrium only exists if n is small and δ and π are large. For this reason we will subsequently restrict attention to the equilibrium described in Proposition 1.4 which exists for all parameter values.

Whether there will be more experiments in the star or the complete network depends on the possibility of the peripheral agents to counterbalance the decreased experimenta-

⁹This terminology is taken from BK, who also refer to equilibria in the star network where some agents contribute to the public good and others do not contribute as specialized.

tion intensity of the hub. In the equilibrium described in Proposition 1.4 the experimentation effort of the hub is below or equal to the effort level of the peripheral players, that is $\phi_t^h \leq \phi_t^s$. For beliefs close to \bar{p}^c (i.e., $\bar{p}^c - \varepsilon$) the agents in the complete network experiment almost with full intensity while hub does not experiment and the agents in the star network cannot increase their effort any more. Consequently, for beliefs right below \bar{p}^c , total experimentation effort is higher in the complete network. The interesting interval are beliefs at which agents in both networks invest in both arms simultaneously. As will be shown in Proposition 1.5, except for a combination of parameter values where n is small and δ and π are large, overall experimentation intensities in the star network are higher or equal to experimentation effort in the complete network in this interval of beliefs.

Proposition 1.5. *Comparing effort exerted in the symmetric equilibrium in the complete network and in the equilibrium in the star network of Proposition 1.4 we obtain*

$$(n-1)\phi_1^s + \phi_1^h = n\phi_1^c \text{ for all } p_1 \in [0, p^a] \cup [\bar{p}^c, 1]$$

and

$$(n-1)\phi_t^s + \phi_t^h = n\phi_t^c \text{ for all } p_t \in [0, \underline{p}_2^s] \cup [\bar{p}^c, 1]$$

where

$$\underline{p}_2^s = \frac{\underline{p}_2^s(1-\pi)^{n-2}}{\underline{p}_2^s(1-\pi)^{n-2} + 1 - \underline{p}_2^s} < p^a.$$

For $p_1 \in (p^a, \bar{p}_1^s)$ there exists a strict subset $S_n(p_1)$ of $[0, 1]^2$ such that

$$(n-1)\phi_1^s + \phi_1^h > n\phi_1^c \text{ if and only if } (\delta, \pi) \in S_n(p_1).$$

Moreover, $\lambda(S_n(p_1)) \rightarrow 1$ as $n \rightarrow \infty$ with λ denoting the Lebesgue measure on \mathbb{R}^2 . Similarly, for $t \geq 2$ and $p_t \in (\underline{p}_2^s, \bar{p}_2^s)$, where

$$\bar{p}_2^s = \frac{\bar{p}_2^s(1-\pi)^{n-2}}{\bar{p}_2^s(1-\pi)^{n-2} + 1 - \bar{p}_2^s},$$

there exists a strict subset $S_n(p_t)$ of $[0, 1]^2$ such that

$$(n-1)\phi_t^s + \phi_t^h > n\phi_t^c \text{ if and only if } (\delta, \pi) \in S_n(p_t)$$

and $\lambda(S_n(p_t)) \rightarrow 1$ as $n \rightarrow \infty$.

The first part of Proposition 1.5 states the intervals of beliefs in which experimentation effort in equilibrium in the complete network is equal to the star network, because there is either no experimentation or all agents exert full effort. For beliefs outside these intervals we know that $\phi_t^h = 0$. The region $S_n(p_1)$ is defined as all combinations of δ and π for which total effort in $t = 1$ in the complete network is strictly smaller than in the star network for $p_1 \in (p^a, \bar{p}_1^s)$. By analyzing this expression (see Appendix A) numerically, one can see that the value for δ below which $(n - 1)\phi_1^s \geq n\phi_1^c$ is in general "quite close" to 1. For example, for $n = 3$, $(n - 1)\phi_1^s \geq n\phi_1^c$ as long as $\delta \leq \frac{8}{9}$ even if π takes values arbitrarily close to 1. As n increases, the threshold value for δ increases and already for relatively small n ($n = 6$) $\delta \leq 0.99$ suffices to guarantee that $(n - 1)\phi_1^s \geq n\phi_1^c$ again assuming values of π close to 1. The lower π , the higher is δ below which $(n - 1)\phi_1^s \geq n\phi_1^c$.

The total experimentation intensity in the interval where agents use both arms in the star network, is higher in the star network except for a combination of parameter values with high δ , high π and small n . That is, unless agents are very patient, effort in the star is higher even though the hub does not experiment. This indicates that the peripheral agents increase own efforts accordingly to outweigh the missing experimentation of the hub as well as the payoff disadvantage that arises from delayed information transmission.

We now turn to a comparison of experimentation intensities in the ring and the star network. As the number of agents increases, more information arrives with a greater number of time lags in the ring. In the star network on the other hand, the delay does not change if the number of players changes. Restricting attention to intervals of beliefs in which neither $\phi_t^r = \phi_t^s = 0$ nor $\phi_t^r = \phi_t^s = 1$, we show in Proposition 1.6 that for small n , experimentation intensities in the star network are no smaller than those in the ring while for a large number of players it depends on δ and π .

Proposition 1.6. *Comparing ϕ_t^s to ϕ_t^r for all p_t from the interval in which at least in one of the two networks agents are indifferent between the safe and the risky option, we have that*

- (i) *there exists $n_t \in \mathbb{N}$ such that for all $n < n_t$, $\phi_t^s \geq \phi_t^r$ for all $(\delta, \pi) \in [0, 1]^2$ and*
- (ii) *as $n \rightarrow \infty$ the region of (δ, π) in which $\phi_t^s \geq \phi_t^r$ is a strict subset of $[0, 1]^2$.*




Figure 1.7: Equilibrium experimentation effort in $t = 1$ in the ring network compared to the star network for $n \rightarrow \infty$. In the light region $\phi_1^s > \phi_1^r$ and in the dark region $\phi_1^s < \phi_1^r$.

The first point of Proposition 1.6 tells us that for a small number of players effort in the star is higher (in the interval of beliefs where agents use both arms) or equal to effort in the ring. For a larger number of agents this is no longer true in general. Part (ii) of the proposition says that as n becomes large, the set of (δ, π) for which effort in the star network is higher becomes a strict subset of $[0, 1]^2$. As a consequence of part (ii) and the fact that ϕ_t^r and ϕ_t^s intersect only at one belief (e.g., in $t = 1$ at p^a), we can conclude that there exists some finite natural number such that for all n above this number there exists a non-empty set of parameters (δ, π) for which $\phi_t^s < \phi_t^r$.

Proposition 1.6 shows that the effect of incomplete network structures on experimentation intensities depends on the discount factor δ and the success rate π . Suppose in both networks all agents experiment in $t = 1$. Then, in $t = 1$ peripheral agents in the star network learn about one experiment fewer than agents in the ring network. For δ close to 1 this does not matter to these agents as it makes little difference to them at which point in time information arrives. On the other hand, the closer δ is to zero, the more agents in the star network care about this one experiment, making them increase own effort. Figure 1.7 compares ϕ_1^s and ϕ_1^r for $n \rightarrow \infty$. In the light region $\phi_1^s > \phi_1^r$ and vice versa in the dark region. For example, $\phi_1^r > \phi_1^s$ only if π is not too large. If the probability of a breakthrough is low, it is relatively more likely that agents will learn about a breakthrough later in the ring than in the star network. Thus, agents in the ring increase their effort to balance this effect.

Before turning to the question which network generates the highest welfare among the three structures considered, let us briefly repeat the main findings of the previous sections. First, we showed that agents increase own effort if information arrives with delay as they are better off experimenting themselves instead of waiting for information generated by others. Second, in irregular structures there can be specialization over some intervals of beliefs where some agents experiment while others free-ride. Experimenting agents increase their effort to outweigh the missing experiments as well as the delay in the information transmission.

1.6 Welfare analysis

In the preceding sections it was shown that effort exerted in equilibrium varies with the network structure. In this section we want to analyze the implications of these differences for expected payoffs in equilibrium.

1.6.1 The optimal network

Assuming that it is costly to establish a communication or interaction structure, we are now interested in which of the three networks would be chosen (before the agents engage in the experimentation game) by a social planner who aims to maximize welfare given the strategic behavior of the players. There are fixed costs $k \geq 0$ per link that have to be paid ex ante. The total number of links in network g depends on the network structure and is $n(n-1)/2$ in the complete network, n in the ring and $n-1$ in the star network. The main criterion to measure the performance of different structures are equilibrium payoffs. Welfare is defined as the total expected payoff in equilibrium minus total costs for building the infrastructure. For the complete network this is

$$W^c(p_1) = nU^c(p_1) - \frac{n(n-1)}{2}k.$$

For the other networks it is defined in an analogous way, that is $W^r(p_1) = nU^r(p_1) - nk$ and $W^s(p_1) = (n-1)U^s(p_1) + U^h(p_1) - (n-1)k$. A network $g \in \{c, r, s\}$ is optimal for a given prior belief p_1 and set of parameters (δ, π, k, n) if and only if

$$W^g(p_1) \geq W^{g'}(p_1), \text{ for all } g' \in \{c, r, s\}.$$

We write $g \succ g'$ if network g generates strictly higher welfare than network g' and $g \sim g'$ if $W^g(p_1) = W^{g'}(p_1)$.

Proposition 1.7 below states which network is optimal when $k = 0$ and the prior belief p_1 is such that in case all experiments in $t = 1$ fail, there are no experiments in $t = 2$, that is, $p_2 \leq \underline{p}_2^g$ for all g . For simplicity of exposition, in the subsequent analysis we impose $E_1 = 1$ and $E_0 = -1$. Note that we do not include the empty network in our analysis, which would be optimal for very high costs. Without the empty network, the star network is optimal for sufficiently high costs. Moreover, the star network is strictly optimal for a certain interval of priors even if links do not incur any costs.

Proposition 1.7. *The following conditions determine which network is optimal for $k = 0$ and for p_1 such that in case all experiments in $t = 1$ fail, $p_2 \leq \underline{p}_2^g$ for all g :*

(i) For $p_1 \in [0, p^a] : c \sim r \sim s$;

(ii) for $p_1 \in (p^a, \bar{p}_1^s] : s \succ c, r$ and the relation between c and r is given in (iii);

(iii) $c \sim r$ for $p_1 \in [0, \bar{p}_1^r]$ and $c \succ r$ for $p_1 \in (\bar{p}_1^r, 1]$;

(iv) for $p_1 \in (\bar{p}_1^s, \bar{p}_1^c] : c \succ s$ if and only if

$$(1 - \delta)(2p_1 - 1) + \delta p_1[(1 - \pi)^{n-1}[1 + \delta(n - 1)] + (1 - \delta)(n - 1)(1 - \pi) - n(1 - \phi_1^c \pi)^n] > 0.$$

(v) For $p_1 \in (\bar{p}_1^c, 1] : c \succ r, s$.

For $p_1 \in (p^a, \bar{p}_1^s]$ the complete network is never optimal even if costs for links are zero. This result is somewhat surprising as one might think that it is optimal to have as many links as possible if they are costless to allow a fast flow of information. However, in this interval of beliefs the star network is strictly optimal for two reasons. First, average expected payoffs in the star (where the hub does not experiment) are higher than in the complete network or the ring, because the hub does not bear the costs of experimentation but receives the informational benefits. Second, up to \bar{p}_1^s the peripheral players can increase their experimentation effort so as to fully compensate for both the lack of experimentation of the hub as well as the delay in the information transmission. Therefore, up to this threshold, therefore, welfare in the star network is strictly higher than in the ring or

the complete network. At some belief above this threshold this result is reversed and the missing experiment of the central player implies that average expected payoffs are lower in the star network than in the other networks. Corollary 1.1 summarizes this result.

Corollary 1.1. *Specialization in the star network, where $\phi_1^h = 0$ and $\phi_1^s > 0$, can be beneficial as well as detrimental to overall welfare.*

Another interesting observation can be made by comparing the complete network to the ring. For beliefs in the interval $(p^a, \bar{p}_1^r]$ agents exert higher effort in the ring network than in the complete network (see Section 1.4). More precisely, agents increase their effort to exactly offset the payoff disadvantage resulting from the delay with which information arrives. This means that expected payoffs in the ring and the complete network are identical for beliefs at which the players in the ring use interior experimentation intensities if there are no costs for links.¹⁰ If all agents in both networks experiment with full intensity agents learn faster in the complete network and are better off. This implies that, as stated in Corollary 1.2, there exists a trade-off between delay and free-riding.

Corollary 1.2. *In the selection of the optimal network structure there exists a trade-off between the speed of learning and the accuracy of learning.*

This trade-off is also apparent when looking at a situation where some agents experiment in $t = 2$ after a round of failed experimentation in $t = 1$. It is possible that in equilibrium in $t = 2$ only the peripheral players in the star network experiment. One main advantage of the complete network compared to incomplete structures lies in the speed of learning, making it increasingly attractive the stronger future payoffs are discounted. In the interval of beliefs in which only the peripheral players experiment in $t = 2$, whereas agents in other networks experiment only in $t = 1$, it can be shown that for values of the discount factor δ close to 1, the star network is always optimal. On the other hand, for δ close to 0, the complete network is optimal for $k = 0$. This comparison stresses again the existing trade-off between learning faster and more precisely, or put differently, between delay and free-riding. How this trade-off is resolved depends on the discount factor.

¹⁰For $k > 0$ the ring network is strictly optimal in the interval $(p^a, \bar{p}_1^r]$.




Figure 1.8: Optimal networks for $k = 0$, $n = 4$ and p_1 such that $p_2 \leq \underline{p}_2^g$ for all g if all experiments in $t = 1$ fail, for different intervals of the prior. The dotted line between \bar{p}_1^s and \bar{p}_1^c indicates the belief at which peripheral players in the star network can no longer compensate for the missing experiment of the hub and the delay in information transmission.

Whether a certain network is optimal, depends on the agents' possibility to increase their experimentation effort in order to compensate for the disadvantage of delayed information arrival in incomplete structures. For costs of links equal to zero, the complete network can only be optimal for prior beliefs p_1 such that $\phi_1^s = \phi_1^r = 1$, as otherwise agents can increase their experimentation effort in order to outweigh the delayed arrival of information. In the star network an additional effect comes into play, namely the payoff advantage of the non-experimenting hub, which explains why even for zero costs the star is strictly preferred for low priors. At some belief in the interval $(\bar{p}_1^s, \bar{p}_1^c]$ the peripheral players in the star can no longer compensate for the nonexperimenting hub and total experimentation effort is lower than optimal. Figure 1.8 graphically illustrates for $n = 4$ which of the three networks is optimal on different intervals of priors.

A fast flow of information does not necessarily maximize welfare even if information can be distributed to all players immediately at no cost due to the strong incentive to free-ride. This contradicts the findings of Teece (1994) that innovation has to be associated with a fast transmission of information. Information is a public good and the underprovision of this public good can be ameliorated by delaying the arrival of information. Our analysis confirms two results of BK. First, it shows that under certain circumstances specialization (that is, some agents exert effort while others free-ride) might benefit society, and second, welfare can be higher in incomplete interaction structures. However, we can also show the opposite effect, namely that for certain beliefs specialization can have a negative impact on overall welfare.

Network structures can also be interpreted as organizational structures that determine the flow of information within an organization. When deciding on the optimal organizational structure (for example, centralized vs. decentralized structures), decision-makers might pursue various objectives. For instance, if the objective is to minimize the costs of information transmission, a centralized structure such as the star network is optimal. Centralization enables a comparatively fast flow of information at lowest possible costs. From the perspective of the management of a firm centralization additionally offers the advantage that a central authority can accumulate and disseminate information.

1.6.2 The complete network with a specialist

The star network is strictly optimal for low prior beliefs in the interval where the hub does not experiment, i.e., in the interval where the asymmetry of the network structure leads to an asymmetry of actions. Asymmetric equilibria in the complete (or ring) network will most likely generate higher welfare than the symmetric one (see KRC or Bramoullè, Kranton, D'Amours, 2014). Consequently, it is not clear whether the star is still strictly preferred once the restriction to symmetric actions in the complete network or the ring is relaxed. In this section we want to find out whether the star network is still optimal in case we allow for some asymmetry in the complete network. There are potentially many asymmetric equilibria and a thorough characterization of these equilibria is in general difficult. Hence, instead of focusing on asymmetric equilibria directly we introduce asymmetry by exogenously imposing specialization: One agent in the complete network, who will be referred to as the "specialist", never exerts any effort and this is commonly known. All other agents choose the optimal experimentation effort prescribed by the symmetric equilibrium of the experimentation game given the specialist.

For beliefs where the hub does not experiment, expected utility in the star network is strictly higher than in the complete network with one specialist. To be more precise, the expected payoff of the peripheral players in the star network is equal to the expected payoff of the working agents in the complete network. Hence, any difference in payoffs results from the difference between the hub and the specialist. Both of them have the same number of direct neighbors, $n - 1$, and can observe them directly. The peripheral agents, however, exert higher effort than agents in the complete network to counterbal-

ance the delay with that information arrives. Thus, the probability of a success in the star network is higher and as a consequence the hub obtains higher expected payoffs than the specialist.

Proposition 1.8. *For $p_1 \in (p^a, \bar{p}_1^s]$ expected payoffs in the star network are strictly higher than expected payoffs in the symmetric equilibrium in the complete network with one non-working specialist.*


Proposition 1.8 implies that the star network is optimal for pessimistic priors even if we allow for some asymmetry in the complete network. Thus, the star network does not only generate higher welfare than the complete network because the network structure leads to specialization, but because of a combination of specialization and delay.

1.6.3 Numerical example


In this section we present numerical results that complement the preceding analytical discussion. While up to this point we focused on the role of the prior, we now want to obtain a better understanding of the role of different parameters. In our numerical example we show which network is optimal in a (π, δ) -grid given fixed values of the other parameters.

Figure 1.9 illustrates the results. It shows which network is optimal for $E_0 = -1$ and $E_1 = 1$ if agents can only experiment in $t = 1, 2$. The results are calculated for $\pi \in [0.01, 0.99]$ and $\delta \in [0.01, 0.99]$ both in steps of 0.01. In white areas all networks are optimal as in this region there is no experimentation (that is, we have indifference). Light gray areas indicate all combinations of δ and π in which the ring network is optimal, in dark gray the star network is optimal, and black means that the complete network is optimal. The three panels on the left display the results for $n = 5$, while those on the right have $n = 25$. In the first row $p_1 = 0.45$ and $k = 0$, in the second row the prior belief is increased to $p_1 = 0.96$ while $k = 0$, and in the last row we look at $p_1 = 0.96$ for costs $k = 0.001$.


In Figure 1.9a we see that for low values of δ and π no network is strictly optimal, as no agent experiments. For medium values of δ , e.g., $\delta = 0.4$, we have indifference




(a) $p_1 = 0.45, k = 0, n = 5$.



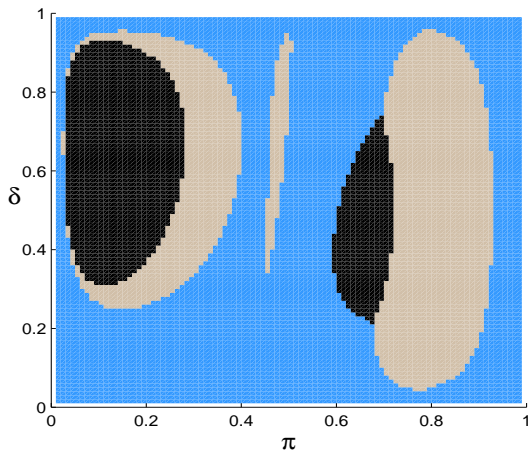
(b) $p_1 = 0.45, k = 0, n = 25$.



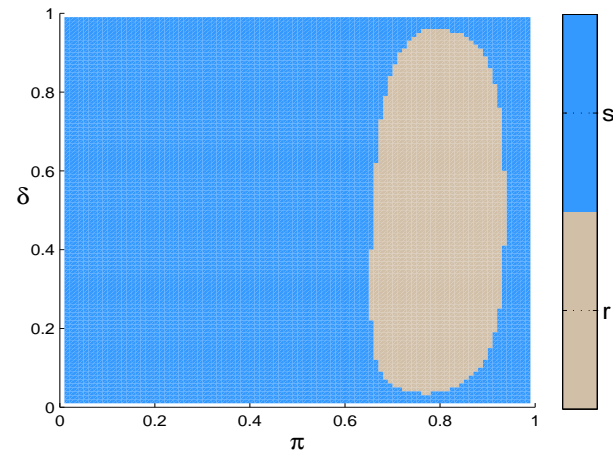
(c) $p_1 = 0.96, k = 0, n = 5$.



(d) $p_1 = 0.96, k = 0, n = 25$.



(e) $p_1 = 0.96, k = 0.001, n = 5$.



(f) $p_1 = 0.96, k = 0.001, n = 25$.

Figure 1.9: Optimal networks for $E_0 = -1$ and $E_1 = 1$.

for low values of π and the star network dominates for high π . As δ and π increase, expected welfare is highest in the complete network. More precisely, in 1.9a the complete network is optimal in 37.4% of the cases, the star in 17.6%, the ring network in 7.4% and in 37.6% of the cases we have indifference. If we increase n to 25 (see Figure 1.9b) the ring network is never optimal and the complete network is optimal for values of the parameters where for $n = 5$ the star is optimal. The percentages change to 61.2% for the complete network, 1.3% for the star and 37.5% for indifference. This implies that although agents face better opportunities to free-ride on the experimentation of others in the complete network, welfare is higher than in the incomplete networks.

In Figures 1.9c and 1.9d agents are very optimistic and experiment with certainty. That is, there is no region of indifference. As expected, the complete network is optimal in this case for a large combination of parameters (92.4% in 1.9c and 97.6% in 1.9d). However, for intermediate values of π there exists an area in which the star network or the ring generate higher welfare. Increasing the number of players to $n = 25$ shifts the region in which the star network dominates to the left, that is, to lower values of π . Moreover, the ring network is never optimal as there is too much delay.

In the last row in Figures 1.9e and 1.9f we introduce positive costs for links. Naturally, the region in which the complete network is optimal shrinks for $n = 5$ and completely disappears for $n = 25$. In fact, for $n = 25$ and $k \geq 0.001$ the complete network is suboptimal for all δ, π , and p_1 .

1.7 Discussion

In this chapter we analyzed a dynamic game of strategic experimentation in three different network structures, the complete network, the ring network and the star network. Relative to the complete network agents in the ring network increase their experimentation effort to balance the payoff disadvantage resulting from the delay in information transmission. In the star network there is an equilibrium where the hub experiments with full intensity up to a threshold belief and then stops completely. Although the peripheral players increase their effort relative to the complete network in the interval where the

hub stops "too early", for some beliefs they are not able to fully compensate for the non-experimenting hub. Hence, depending on the belief this specialization where the hub does not experiment can be beneficial as well as detrimental for society.

The obtained results offer insights into the incentives that drive the behavior of rational agents. Taking research or innovation as examples, a welfare analysis of the model provides insights which might be relevant to government authorities or companies for structuring and subsidizing research projects. Objectives of decision makers can be manifold as for instance, cost minimization, utility maximization, the maximization of the speed of learning through fast information transmission or completeness (that is, more precise learning which implies that the probability of mistakenly abandoning a good risky arm is minimized). Different network structures have different effects on the outcome of the experimentation game and consequently on welfare. While the star network minimizes the costs for links, the complete network maximizes the speed of learning. In which of the three networks learning will be most accurate depends on the prior belief as well as the parameters of the model. In general, there exists a trade-off between faster learning and more accurate learning.

Our analysis showed that it is possible to investigate details of rational learning processes in a network without being restricted to focus on asymptotic results or introduce some form of myopia or bounded rationality. Nevertheless, the model considered here captures very particular learning situations due to its special structure with fully revealing breakthroughs. This implies that our results cannot be easily generalized to other payoff generating processes. Another shortcoming of the analysis is the restriction to symmetric equilibria which may not be innocuous. However, comparing expected payoffs in the star network to payoffs in the complete network with one non-working "specialist" showed that allowing for some degree of asymmetry does not change the basic intuition of our results. In incomplete network structures the delay mitigates free-riding and increases the probability of success, which may increase welfare. An incomplete network generates strictly higher payoffs than the complete network only in case both forces, delay and specialization, are present.

Despite the complexity network structures can create, we showed that they affect the

behavior of agents in an intuitive way. This offers some suggestions as to how equilibrium outcomes and strategies could be characterized in other settings (e.g., for other payoff generating processes) as well. Further, the network structures considered in this chapter can be understood as specific monitoring structures, and it would be possible to analyze the strategic experimentation game for monitoring structures which are not derived from networks. What should be clear, however, is that the empty and the complete network are two opposite ends of the spectrum. Consequently, for symmetric monitoring structures we expect the main conclusions of the network case to remain valid. Of course, it would be desirable to obtain a generalization of the results for irregular structures as well, which seems to be considerably more involved and will most likely imply specialization as in the star network. The second equilibrium in the star network in which the central player uses a strategy that is non-monotonic in the belief already points at the potential complexity of equilibria in more complex network structures.

Chapter 2

Innovation, intersectoral R&D spillovers and intellectual property rights¹

2.1 Introduction

The importance or profitability of a new idea or innovation often depends on how innovations in different areas complement each other. For example, the effectiveness of a new drug in the treatment of a certain disease might depend on the availability of machines or tests to diagnose the disease at an early stage. The diffusion of new software or applications depends on the ability of existing hardware to operate such programs. Production of a new product might be facilitated if new machines are available or progress in resource extraction lowers the costs of production. Park (2004) shows that manufacturing R&D has a substantial intersectoral R&D spillover effect on the productivity growth of the nonmanufacturing sector. Consequently, the profitability of innovation activities in one field is linked to the success of developments in other areas and these intersectoral R&D spillovers are the subject of our research.

In this chapter we offer a model to analyze situations in which the aforementioned complementarities in the innovation process are present. We consider a situation where firms operate in one of two sectors (for example, software and hardware). Each firm can engage in research in its respective field of specialization where the profit that can

¹This chapter originated from a joint research project in collaboration with Michael Metzger on complementarities in models of strategic experimentation.

be earned from an invention depends not only on own effort or R&D investment, but also on innovation efforts and outcomes in the other field. More precisely, we analyze a game of strategic experimentation with externalities between different sectors. Our model is based on the discrete time version of the exponential bandit model by Keller et al. (2005) as offered in Heidhues et al. (2015). In games of strategic experimentation² firms commonly face identical two-armed bandit machines where the outcomes of the arms are uncertain. Firms learn about the type of the arms by playing repeatedly and observing the outcomes. This active learning is also called experimentation. In our interpretation firms can invest in research projects of unknown quality which is either good or bad. A bad project never leads to an invention and it is optimal not to invest, whereas investment in a good project leads to a discovery with positive probability and investment is profitable. Such a setup describes situations in which firms invest in research where they do not know whether the research will be successful at all and if so, when they will observe a success. Besanko and Wu (2010) refer to this as research projects facing "if" and "when" uncertainty.

In what follows we analyze the innovation processes in two sectors where innovation activities within one sector are considered as substitutes as in Keller et al. (2005). Innovation activities in one sector, however, are complements to innovation activities in the other sector. To be more specific, we assume that the payoffs in the first sector are independent of what happens in the second sector, while payoffs that can be earned from an invention in the second sector are higher if there was a discovery in the first sector. Hence, we will refer to the first sector as the independent sector and to the second sector as the dependent sector.³

Economic policies that influence the R&D activities of firms can have unexpected consequences if the policy maker is not aware of spillovers across sectors. For instance, the

²See Bergemann and Välimäki (2008) for a survey on multi-armed bandit models and their application in economics.

³Given the findings of Park (2004) that manufacturing R&D has an intersectoral spillover effect on non-manufacturing productivity growth while nonmanufacturing R&D does not have the same effect on manufacturing productivity, we find it plausible to assume an asymmetric relationship between innovations in the two sectors. This allows us to compare research projects that have a complementary character to research projects without this feature.

Federal Trade Commission (2011) [FTC] states that especially in the IT sector, patent notification is characterized by a lack of transparency which might cause firms to postpone investing. At the same time the U.S. National Science Foundation invests into ICT research emphasizing the importance of this research for other fields.⁴ Hence, if policies are implemented that increase transparency in the ICT sector, this affects other research areas as well. In this chapter it is shown what effects one can expect to observe in the presence of spillovers across sectors, which seems to be particularly relevant given the inter- and multidisciplinary character of present day research. More precisely, besides analyzing the effects of a change in intellectual property rights [IPR] on investment within a sector, we as well study the impact on investment in a complementary sector. In particular, we want to know how project choice (that is, which sector a firm selects) is influenced by intersectoral R&D spillovers, and how IPR that guarantee the optimal choice for independent sectors differ from those for dependent sectors.

To answer these questions we analyze a dynamic two-firm model of R&D investment with inter- and intrasectoral R&D spillovers. Our first finding is that whether a firm in the dependent sector invests more in research in the presence of intersectoral spillovers depends on how fast firms reach a belief about the viability of the research at which investment is no longer profitable. If the firm in the independent sector learns fast, it stops investing earlier after a history of failures. This implies that a firm in the dependent sector knows sooner that it will not be able to obtain the high profit that results from two complementary inventions. The firm then invests the same amount as without complementarities. On the other hand, if the firm in the independent sector makes a discovery or learns comparatively slowly, research investment in the dependent sector in the presence of complementarities exceeds investment without complementarities.

Similarly to Besanko and Wu (2010, 2013) we allow for different levels of intrasectoral spillovers. Without intrasectoral spillovers only the firm making the discovery can profit from it, while positive spillovers allow the unsuccessful firm to copy an invention and

⁴See CIF21 (Cyberinfrastructure Framework for 21st Century Science and Engineering). Also in the UK, the (2010) report of the Department of Business, Innovation and Skills on the allocation of research funding shows that research funds are dedicated to bring together ICT research with other areas (as e.g., medicine) to enable multidisciplinary research.

generate profits as well. The easier it is to imitate the invention of another firm (i.e., the higher the level of intrasectoral spillovers), the more firms should invest in R&D from the perspective of a welfare-maximizing social planner. However, R&D investment of strategic firms decreases in the level of intrasectoral spillovers, because the easier it becomes to imitate, the stronger is the incentive of firms to reduce own R&D investment and free-ride on the R&D investment of the other firm.

After analyzing R&D investment for a given choice of sectors or research lines, we concentrate on the interplay between inter- and intrasectoral R&D spillovers when firms select the line of research in the beginning of the game. If welfare does not depend on the level of intrasectoral R&D spillovers, a social planner prefers joint research in the independent sector (to increase the probability of a discovery) for beliefs where strategic firms pursue different lines of research. Whether the social planner prefers joint research in the dependent sector depends on the level of intersectoral spillovers. More precisely, the social planner faces a trade-off between increasing the probability of success in the dependent sector through joint research, and having firms work in different sectors which allows the firm in the dependent sector to obtain a higher payoff after two complementary discoveries. Which effect dominates depends on the extra payoff associated with two complementary innovations and the probability of a discovery. If welfare is increasing in intrasectoral spillovers, the social planner prefers diversification if firms face a winner-takes-all competition to avoid too much competition on one line. For perfect positive spillovers the social planner prefers joint research in the independent sector, while joint research in the dependent sector is only preferred if intersectoral spillovers are low.

Finally, we ask how firms can be encouraged to choose the socially optimal research line. We assume that this decision can be influenced through the design of IPR, which are interpreted as policies that influence the level of intrasectoral spillovers. If consumer surplus is decreasing in the level of intrasectoral spillovers, imitating a successful invention in the dependent sector is exacerbated to encourage an entering firm to locate in the independent sector and thereby create positive intersectoral R&D spillovers. If welfare is increasing in intrasectoral spillovers, imitating an invention in the dependent sector is facilitated to foster joint research.

This chapter is the first to study intersectoral R&D spillovers in a model of innovation with uncertainty. While Steurs (1995) studies intersectoral spillovers without uncertainty, Besanko and Wu (2010, 2013) focus on the impact of uncertainty on R&D investment, if research is concentrated in one sector only. The combination of intersectoral R&D spillovers and uncertainty determines not only R&D investment within a sector but also the selection of research lines. Knowing how such spillovers affect a firm's decision to invest in a certain research line offers insights into how intellectual property rights interact with intersectoral spillovers and how they can be used to influence the firms' investment decisions.

This chapter is structured as follows. Section 2.2 introduces the model. In Section 2.3 we analyze strategic and socially optimal R&D investment after firms have selected a sector. After that, in Section 2.4, we analyze the decision of selecting a sector in the beginning of the game, before discussing IPR in Section 2.5. Section 2.6 contains a discussion and conclusion. All proofs are relegated to Appendix B.

2.1.1 Related literature

Intersectoral R&D spillovers have been studied by Steurs (1995) in a two-industry, two firms-per-industry model. Intersectoral spillovers affect R&D investment directly and indirectly through changing the influence of intrasectoral spillovers. Further, research cooperatives across sectors might be more socially beneficial than cooperatives in one sector. Unlike in the model considered subsequently R&D in Steurs (1995) is not associated with uncertainty regarding feasibility or timing. That is, research leads to a cost reduction with certainty. Moreover, while Steurs (1995) focuses on the impact of intersectoral spillovers on output and profits, we concentrate on the selection of research lines and investigate how intrasectoral spillovers (or IPR) can be used to encourage firms to select optimally between research lines.

A few papers consider complementarities in innovation and the consequences of complementarities for the design of intellectual property rights (see e.g., Gancia and Zilibotti, 2005 or Young, 1993). More recently, Chen (2012) analyzes the innovation frequency of durable goods that are perfect complements and shows that interdependencies of inno-

vation decisions lead to coordination failures between producers. While Chen (2012) focuses on product complementarity Bessen and Maskin (2009) consider technological complementarities. More precisely, they show that when innovation is sequential (each successive innovation builds on its predecessors) and complementary (each innovator takes a different research line), patent protection is not necessarily useful for encouraging innovation. Hunt (2004) also studies sequential innovation in a model with endogenous industry structure and shows that there is a unique patentability standard that maximizes the rate of innovation. Our model differs from these papers in the following ways. First, we explicitly assume that, while inventions might complement each other, a success in one sector is not necessary to generate profits from innovations in other fields, so we do not have perfect complements. Second, we do not consider a sequential setting in which innovations build on each other. That is, a success in one sector does not affect the probability of success in the other sector.

Amir et al. (2003) analyze cooperative and non-cooperative R&D investment and characterize the profit-maximizing R&D cartel when both R&D investment and intrasectoral spillovers are cooperatively chosen. In the profit-maximizing cartel the spillover is either maximal and firms cooperatively choose their R&D investment or only one firm invests and the spillover is minimal. Miyagiwa and Ohno (2002) analyze the impact of uncertainty on cooperative R&D in the presence of "when"-uncertainty and show that firms want to coordinate R&D investment but not necessarily share the innovation.

Besides this, the chapter relates to the literatures studying the role of competition in models of strategic experimentation. Akcigit and Liu (2015) investigate the effect of competition if firms can choose between a risky and a safe research line and do not observe the actions and failures of others. Two types of inefficiencies arise; one due to firms switching to the safe alternative too early and one due to wasteful duplication of R&D effort. Besanko and Wu (2010, 2013) and Das (2013) study investment in an exponential bandit framework with payoff externalities. Acemoglu et al. (2011) show that patents improve the allocation of resources in a model of experimentation. These models differ from ours in that they consider only one sector and there are no intersectoral R&D spillovers or complementarities between different lines of research.

2.1.2 Examples

Firms can invest in uncertain research in different fields or sectors. For instance, the development of a new drug in the pharmaceutical industry or medical sciences, and new techniques for data analysis in the field of computer sciences or statistics. In general, it is not clear whether a success can be achieved or when it will be achieved. While firms invest in research (e.g., on the drug), they become more and more pessimistic that the development of a new drug is feasible if they do not achieve a breakthrough. After the first discovery both firms know that this research line is viable and, for example, a cure of the disease is possible. If one firm conducts research on a new drug while the other develops new tools for data analysis, a success of either one does not tell the other firm whether it will be successful as well. New techniques for data analysis or advances in ICT might, however, facilitate the analysis of clinical trials and thereby lower the costs of development for the drug. This example incorporates several features of our model. First, the discovery of a new drug does not affect the profitability of new data analysis techniques. Second, even without progress in data analysis a drug that was found to be effective can be profitably sold, and third, improvements in data analysis alone do not generate any profit for the pharmaceutical firm if it cannot develop the drug.

There are several empirical examples of inventions that feature intersectoral R&D spillovers. For instance, researchers at the Mayo clinic conducted a study on gene expression in the brain to improve understanding of Alzheimer's disease. The researchers used genetic interaction studies in which effects of pairs of gene changes are studied. This process involves the analysis of billions of DNA base pairs and requires substantial computational processing time. By using the Blue Waters supercomputer the computation time was reduced from more than a year to merely two days.⁵

2.2 Model

The model is based on the discrete time strategic experimentation model of Heidhues et al. (2015) with the main difference that there are two sectors (research lines, projects)

⁵See U.S. National Science Foundation (2015).

$s = 1, 2$. There are two firms $i = 1, 2$ that can invest in R&D in discrete time in $t = 1, 2, \dots, T$ and discount future payoffs by a common discount factor $\delta \in (0, 1)$.

2.2.1 Timing

The game proceeds in three stages.

Stage 1: Selecting a sector

In the first stage of the game, in $t = 0$, firms decide which research line to pursue, i.e., they decide in which of the two sectors to locate. This decision is irreversible and determines in which of the two research lines a firm can invest in stage 2. For example, high sunk costs to start research on a particular research line might make a later change to a different field too costly.

Stage 2: R&D investment

In the second stage, the research stage, in each sector $s = 1, 2$ firms can invest up to one unit of an available resource (e.g., money or effort) in research in each period of time until time $T < \infty$. The action or investment decision at time t is denoted by $k_{s,t} \in [0, 1]$.⁶ If $k_{s,t} = 0$ the firm does not invest in research, and $k_{s,t} = 1$ means that the firm invests its entire resource. Firms are not allowed to experiment after time T , that is, $k_{s,t}$ is restricted to be 0 for all $t > T$.⁷ R&D investment is costly and costs $c > 0$ are proportional to investment meaning that a firm choosing investment level $k_{s,t}$ pays costs $k_{s,t}c$ at time t .

R&D investment can lead to a discovery depending on the state of the world. The state of the world in sector 1 is independent of the state in sector 2 and for each sector it can be either good or bad. If it is bad no discovery can be made and benefits of research are zero. If it is good discoveries occur with probability $k_{s,t}\lambda_s$. Discoveries in sector 1 and sector 2 are independent of each other and independent across time and across players. Upon the first discovery (also called breakthrough) firms know that a sector is good. A

⁶Clearly $k_{s,t}$ depends on i as well, but we do not make this explicit in the notation (unless necessary) to enhance readability.

⁷This assumption allows us to solve for the optimal strategy using backward induction. Moreover, in many instances research funds are granted for a limited period of time and are only prolonged in the case of a success, so that arbitrarily long experimentation without success is not possible.

discovery is the beginning of the third stage in which the invention is sold in the product market. That is, a discovery in period t generates profits from period $t + 1$ onwards.

Stage 3: Product market competition

The profit firms earn from an invention depends on the situation in the product market and we distinguish two scenarios. First, each firm can be in a different sector and hence acts as a monopolist in this sector. Second, firms compete in the product market in the same sector.

Upon discovery a monopolist in sector $s = 1, 2$ sells one unit each period at zero marginal costs and price x_s . We assume that the monopolist is able to reap the entire consumer surplus and hence the price equals the consumers' willingness to pay for the discovery. Payoffs in sector 1 are unaffected by sector 2 and conditional on a successful discovery the monopolist receives a per period profit of x_1 which we normalize to 1. In sector 2 we have the following situation. If sector 2 is good a monopolist obtains \bar{x} in case there was a discovery in sector 1, and \underline{x} in case there was no discovery in sector 1, where we assume that $\bar{x} > \underline{x}$. Subsequently we will refer to sector 1 as the independent sector and to sector 2 as the dependent sector.

In case both firms are in the same sector competition in the product market affects profits and consumer surplus. We assume that after a discovery every period each firm sells one unit and collusive behavior is prohibited. We denote the consumer surplus in sector s if both firms compete in this sector in the product market by ψ_s and assume $\psi_s > 0$. Further, competition lowers the profit each firm obtains to αx_s , where $\alpha \in (0, 1)$.

2.2.2 Intrasectoral R&D spillovers

If both firms work on the same research line intrasectoral R&D spillovers determine how easily the discovery of one firm can be imitated or used by the unsuccessful firm. In the benchmark models by Keller et al. (2005) or Heidhues et al. (2015), all firms benefit equally from one success so that it does not matter who achieves the breakthrough. While this seems to be a plausible assumption if players are members of a team who share a common goal, we might as well find situations in which the agent who has a breakthrough receives some sort of reward. For example, firms might enjoy some periods of monopoly

profits in which other firms are not allowed or able to copy their innovation, employees might receive a bonus on top of their salary, or players might simply gain utility from the achievement of a discovery itself.

More precisely, we assume that the firm who has the breakthrough receives αx_s , while the other firm receives $\gamma_s \alpha x_s$ in sector $s = 1, 2$, with $\gamma_s \in [0, 1]$. If $\gamma_s = 0$ we have a winner-takes-all race in which only the inventor receives a high profit, while $\gamma_s > 0$ corresponds to a situation in which a discovery has positive spillovers and can be profitably imitated by the unsuccessful firm. In case both firms simultaneously make a discovery, payoffs have to be shared and each firm obtains a share of $\frac{1+\gamma_s}{2}$. Besanko and Wu (2010) use the same way to model intrasectoral spillovers, whereas in Besanko and Wu (2013) negative spillovers are possible because an invention has a negative impact on an established product of the unsuccessful firm.⁸

The parameter γ can also be interpreted as a measure of intellectual property rights. A small γ reflects a situation in which an invention cannot be imitated by a competitor easily. This can be due to long patent periods or strong patent protection in which patents are wide in scope. A high γ corresponds to an environment in which an innovation can be easily copied and profitably used by an imitator, for example, because patent periods are short or patents are narrow so that an imitator can sell a similar product. In particular, inventions might be adapted for sale in a different geographical region or the technique of the inventor might be applied by the imitator in a different context.

2.2.3 Beliefs and informational externalities

Firms attach a probability to each sector being good and we assume that they start with common priors. At time t the probability firms attach to sector $s = 1, 2$ being good is denoted by $p_{s,t}$. We assume that all actions and outcomes are perfectly observable by all players. When talking about firms in different sectors this assumption might seem unrealistic. One can, however, simply consider the problem of a corporation with subsidiaries being active in different branches interpreting player 1 as an employee in one department

⁸Negative spillovers correspond to $\gamma_s < 0$ and are not considered in the subsequent analysis.

or subsidiary and player 2 as an employee in another one. Apart from that, players can also be interpreted as researchers in different institutes or departments at one university.

Beliefs depend on the occurrence of a breakthrough and in a situation where all actions and outcomes are perfectly observable common posterior beliefs prevail. Beliefs are updated according to Bayes' rule meaning that the subjective probability that sector s is good after investing $k_{s,t}$ without a discovery is given by

$$p_{s,t+1} = \frac{p_{s,t}(1 - k_{s,t}\lambda_s)}{p_{s,t}(1 - k_{s,t}\lambda_s) + 1 - p_{s,t}}, \quad (2.1)$$

where $1 - k_{s,t}\lambda_s$ is the probability of experimenting without success in s conditional on s being good. In case of a breakthrough the posterior belief jumps to 1. Every time firms experiment without success, they become more pessimistic about the research being viable. Firms are risk neutral expected profit maximizers and a firm is said to experiment (or invest in R&D) if it invests (chooses $k_{s,t} > 0$) before knowing the state of the world.

In a good sector 2 the expected profit is $x_2 = \underline{x} + (\bar{x} - \underline{x})p_{1,t}q_t$, where q_t equals $k_{1,t}\lambda_1 + \lambda_1 \sum_{r=t+1}^T \delta^{r-t}k_{1,r} \prod_{l=t}^{r-1} (1 - \lambda_1 k_{1,l})$ if there was no breakthrough in sector 1 up to time t , and 1 otherwise. Denoting $\bar{x} - \underline{x}$ by \tilde{x} , $x_2 = \underline{x} + \tilde{x}p_{1,t}q_t$ equals \bar{x} if there was a breakthrough in sector 1, as in this case $q_t = 1$ and $p_{1,t} = 1$. To avoid the trivial case where it is never optimal to invest in R&D we impose

$$\delta\lambda_s x_s \alpha > c(1 - \delta), \quad (2.2)$$

which implies that in a good sector expected profits exceed the costs of research. That is, in each sector it is optimal to invest ($k_{s,t} = 1$) in case the state of the world is good and choose $k_{s,t} = 0$ if the state is bad.

2.2.4 Producers, consumers and welfare

Producer surplus is increasing in γ_s , i.e., producers are better off, the easier it is to imitate the invention of the other firm. Consumer surplus, however, might be decreasing in γ_s . In what follows we will distinguish two possible cases. First, we consider the case where consumer surplus is decreasing in γ_s and welfare does not depend on γ_s . This means that γ_s shifts the available surplus between consumers and producers. Suppose the consumers' valuation of two units is given by $\underline{x}_s(1 + \alpha)$ and the firms generate a profit of

$\underline{x}_s \alpha(1 + \gamma_s)$, which also equals the price that has to be paid by the consumers. Then the consumer surplus ψ_s is given by $\underline{x}_s(1 - \alpha\gamma_s)$ and welfare $W_s = \underline{x}_s(1 + \alpha)$. Second, we analyze the case where consumer surplus does not depend on intrasectoral R&D spillovers and hence welfare is increasing in γ_s . This means that γ_s does not simply shift the surplus between consumers and producers, but determines the size of the surplus. Subsequently we will focus on the results for $\psi_s = 0$.⁹

2.2.5 Investment strategies and equilibrium concept

Starting from $t = 1$ firms choose their R&D investment $k_{s,t} \in [0, 1]$ given their own location as well as the location of the other firm. An investment strategy for a firm in sector 1 is a mapping that assigns to each belief $p_{1,t}$ a player can have about sector 1 an action $k_{1,t} \in [0, 1]$ for $t = 1, 2, \dots, T$. In sector 2 a pure strategy assigns to each tuple $(p_{1,t}, p_{2,t})$ an action $k_{2,t} \in [0, 1]$ for $t = 1, 2, \dots, T$. We focus on equilibria in Markovian strategies. Moreover, if both firms are located in the same sector they are restricted to symmetric strategies. In the next section we analyze R&D investment for a given choice of sectors, before turning to the selection of a sector in Section 2.4.

2.3 R&D investment with intersectoral R&D spillovers

As a first step we derive the investment strategy for a given sector, that is, we start our analysis in stage 2. We start by analyzing the R&D investment of firms when each firm is located in a different sector in the presence of intersectoral R&D spillovers from the independent sector (e.g., computer sciences) to the dependent sector (e.g., medical research). After that, we look at the situation where both firms research in the same sector. Strategic R&D investment is compared to the welfare-maximizing R&D investment of a social planner who takes the profit of both firms and the consumer side into account.

⁹The results do not change qualitatively if we consider a different payoff specification as long as welfare is increasing in γ_s . For example, suppose the consumers value two units at $\underline{x}(1 + \alpha)$, and the price they have to pay equals $2\alpha\underline{x}$, which means that consumer surplus is given by $\underline{x}(1 - \alpha)$. Moreover, the imitator has to pay costs for imitating the discovery that are decreasing in γ_s and given by $\underline{x}\alpha(1 - \gamma_s)$. Then producer surplus equals $\underline{x}\alpha(1 + \gamma_s)$ and welfare $\underline{x}(1 + \alpha\gamma_s)$.

2.3.1 One firm in each sector

First, we analyze the behavior of firms when one firm is located in each sector, where we assume that firm 1 is located in sector 1 and firm 2 in sector 2.

Strategic R&D investment

Let us start by looking at the independent sector 1. The firm maximizes its total expected profit given by

$$(1 - \delta)E \left[\sum_{t=1}^{\infty} \delta^{t-1} (\delta X_{1,t} - k_{1,t}c) \right] \quad (2.3)$$

by choosing an optimal action profile $\{k_{1,t}\}_{t=1}^T$. The expectation is taken w.r.t. the belief $p_{1,t}$ and the strategy, and $X_{1,t} \in \{0, 1\}$ is the profit resulting from research in sector 1 at time t . By Bellman's Principle of Optimality we can rewrite the profit maximization problem recursively so that the value function of the firm in the independent sector at a given point in time satisfies

$$u_1(p_{1,t}) = \max_{k_{1,t} \in [0,1]} \{-(1 - \delta)k_{1,t}c + \delta E[u_1(p_{1,t+1})|p_{1,t}, k_{1,t}]\}. \quad (2.4)$$

The first term on the r.h.s., $-(1 - \delta)k_{1,t}c$, represents the normalized costs of research. The second term, $\delta E[u_1(p_{1,t+1})|p_{1,t}, k_{1,t}]$, represents the discounted expected continuation payoff. This continuation payoff equals 1 with the probability $p_{1,t}k_{1,t}\lambda_1$ that the sector is good and the firm has a breakthrough at time t . The continuation payoff equals $u_1(p_{1,t+1})$ with the probability $1 - p_{1,t}k_{1,t}\lambda_1$ that there is no breakthrough, where $p_{1,t+1} = \frac{p_{1,t}(1 - \lambda_1 k_{1,t})}{1 - p_{1,t}\lambda_1 k_{1,t}}$. If a firm invests without success its belief declines, and as expected payoffs are increasing in the belief, expected payoffs from investing decrease over time in the absence of a breakthrough. Consequently, the firm invests for high or optimistic beliefs and stops for low or pessimistic beliefs. Lemma 2.1 states the optimal investment strategy for a monopolist in sector 1. The firm follows a time-invariant cutoff strategy with cutoff belief p_1^c . This threshold belief is increasing in costs and decreasing in δ and λ_1 . That is, as firms become more patient or the probability of a success increases, firms invest more in R&D and the likelihood of a discovery increases.

Lemma 2.1. *In the independent sector the monopolist's optimal investment strategy for $t \leq T$ is*

$$k_{1,t}^a = \begin{cases} 1 & \text{for } p_{1,t} \in [p_1^a, 1], \\ 0 & \text{otherwise,} \end{cases}$$

where

$$p_1^a = \frac{(1 - \delta)c}{\delta\lambda_1}.$$

Let us now turn to the dependent sector. The firm's value function satisfies

$$u_2(p_{1,t}, p_{2,t}) = \max_{k_{2,t} \in [0,1]} \left\{ -(1 - \delta)k_{2,t}c + \delta E[u_2(p_{1,t+1}, p_{2,t+1}) | p_{1,t}, p_{2,t}, k_{1,t}, k_{2,t}] \right\}. \quad (2.5)$$

The first term in braces on the r.h.s. again represents the costs of R&D. The continuation payoff is $x_2 = \underline{x} + \tilde{x}p_{1,t}q_t$ with the probability $p_{2,t}\lambda_2k_{2,t}$ that firm 2 observes a success, $u_2(1, p_{2,t+1})$ with the probability $(1 - p_{2,t}\lambda_2k_{2,t})p_{1,t}\lambda_1k_{1,t}$ that firm 2 does not make a discovery but firm 1 does, and $u_2(p_{1,t+1}, p_{2,t+1})$ with the probability $(1 - p_{2,t}\lambda_2k_{2,t})(1 - p_{1,t}\lambda_1k_{1,t})$ that neither makes a discovery. As above, we can derive the belief below which investing is no longer optimal in the dependent sector. This belief depends on the actions and outcomes in the independent sector as can be seen in Lemma 2.2.

Lemma 2.2. *In the dependent sector the monopolist's optimal investment strategy for $t \leq T$ is*

$$k_{2,t}^a = \begin{cases} 1 & \text{for } p_{2,t} \in [p_2^{stop}(p_1), 1], \\ 0 & \text{otherwise,} \end{cases}$$

where

$$p_2^{stop}(p_1) = \frac{(1 - \delta)c}{\delta\lambda_2(\underline{x} + p_{1,t}q_t(p_1)\tilde{x})},$$

and

$$q_t(p_1) = \begin{cases} 1 & \text{for } p_{1,t} = 1, \\ k_{1,t}^a\lambda_1 + \lambda_1 \sum_{r=t+1}^T \delta^{r-t} k_{1,r}^a \prod_{l=t}^{r-1} (1 - \lambda_1 k_{1,l}^a) & \text{otherwise.} \end{cases}$$

If firm 1 stops investing before firm 2, then $q_t(p_1) = 0$ and we denote the autarky cutoff p_2^{stop} at $q_t = 0$ by p_2^a . This is also the belief at which the firm in sector 2 stops investing in case there are no intersectoral spillovers, meaning $\tilde{x} = 0$. How many time periods it takes until the posterior belief falls below the respective cutoff depends, among others, on the success probabilities λ_1 and λ_2 , as they determine how strong the belief declines after observing a failure. Additionally, they specify the value of the stopping cutoff. If $\lambda_1 > \lambda_2$, p_1

declines faster than p_2 , but the stopping cutoff p_1^a is smaller or more pessimistic than p_2^a . If all parameters of the model are the same (prior, discount factor and payoffs), firms with a higher λ might stop at an earlier or later point in time, as it is not clear which effect dominates (see also Halac, Kartiv and Liu, 2013). Following a breakthrough in the independent sector, $q_t = 1$ and $p_{1,t} = 1$ and the firm in the dependent sector experiments until p_2^{stop} at $q_t(p_1) = 1$, which we denote by p_2^b . In case firm 1 still experiments but there was no breakthrough in the independent sector so far, $q_t(p_1) = k_{1,t}^a \lambda_1 + \lambda_1 \sum_{r=t+1}^T \delta^{r-t} k_{1,r}^a \prod_{l=t}^{r-1} (1 - \lambda_1 k_{1,l}^a)$ and the corresponding cutoff is denoted by $p_2'(p_1)$. We have $p_2^a > p_2'(p_1) > p_2^b$. The probability of making a discovery in the independent sector decreases over time and so does q_t , which means that $p_2'(p_1)$ increases and goes to p_2^a as $q_t(p_1)$ decreases.

We are particularly interested in whether intersectoral spillovers induce firms to invest more in research than they would in the absence of the complementarity, that is, whether experimentation beyond the autarky solution is possible. From Lemma 2.2 we see that a breakthrough in the independent sector increases expected profits from a discovery in the dependent sector and causes firm 2 to experiment longer. Experimentation below p_2^a in the dependent sector without a breakthrough in the independent one is, however, only possible if p_1^a is reached after the firm in the dependent sector reaches p_2^a . This can happen if, for example, firm 1 is sufficiently more optimistic in the beginning. Further, it is possible that firm 2 already stopped experimenting, but invests again if firm 1 makes a discovery after firm 2 stopped. Corollary 2.1 summarizes these results.

Corollary 2.1. *A firm in the dependent sector experiments for beliefs below the autarky cutoff p_2^a if and only if one of the following is satisfied:*

- (i) *There is a breakthrough in the independent sector.*
- (ii) *In the absence of a breakthrough, the firm in the independent sector reaches the cutoff belief p_1^a after the firm in the dependent sector reaches p_2^a .*

Welfare-maximizing R&D investment

Suppose now that a welfare-maximizing social planner, who takes the profits of the firms and the benefits to consumers into account, can choose the optimal R&D investment of

both firms. With one firm in each sector consumer surplus is zero and the socially optimal R&D investment coincides with the optimal R&D investment of firms in a research cooperation aiming to maximize joint profits. Expected aggregate payoffs from action profile $(k_{1,t}, k_{2,t})$ after a history of failures are given by

$$W(p_{1,t}, p_{2,t}, k_{1,t}, k_{2,t}) = -(1 - \delta)c(k_{1,t} + k_{2,t}) + \delta p_{1,t} \lambda_1 k_{1,t} + \delta p_{2,t} \lambda_2 k_{2,t} + \delta(1 - p_{1,t} \lambda_1 k_{1,t})u_1(p_{1,t+1}) + \delta(1 - p_{2,t} \lambda_2 k_{2,t})u_2(p_{2,t+1}, p_{2,t+1}).$$

Expected aggregate payoffs $W(p_{1,t}, p_{2,t}, k_{1,t}, k_{2,t})$ are linear in $k_{1,t}$ and $k_{2,t}$. Therefore, the socially optimal strategy can be found by calculating and comparing the aggregate payoff for $(k_{1,t}, k_{2,t}) = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$ assuming there are no experiments in $t + 1$. Denoting W as a function of $(k_{1,t}, k_{2,t})$ only, we see that the same cutoff beliefs as for strategic R&D investment, p_1^a and p_2^a , determine when $W(1, 0) \geq W(0, 0)$ and $W(0, 1) \geq W(0, 0)$ respectively. Furthermore, $W(1, 1) \geq W(1, 0)$ if $p_2 \geq p_2'(p_1)$, which is as well the same cutoff as in the strategic setting. Finally, $W(1, 1) \geq W(0, 1)$ if p_1 is greater than or equal to

$$\frac{(1 - \delta)c}{\delta \lambda_1 (1 + \lambda_2 p_2 \tilde{x})} < p_1^a. \quad (2.6)$$

In the dependent sector strategic and welfare-maximizing R&D investment coincide, that is, firms experiment until the same cutoff belief. In the independent sector, however, the social planner experiments beyond the single firm cutoff, because of the positive effect on expected payoffs in the dependent sector (see also Figure 2.1). A strategic firm in the independent sector does not take the effect of its investment on the other firm into account and stops at too optimistic a belief. This inefficiency can be completely ascribed to the complementarity between sectors.

2.3.2 Two firms in one sector

Let us now analyze R&D investment if both firms are in the same sector. In this case intrasectoral spillovers γ_s indicate how easily the discovery of one firm can be imitated by the other firm and moreover competition in the product market decreases profits to αx_s and $\alpha \gamma_s x_s$ respectively. We assume that also for $\gamma_s = 0$ (which implies a profit of zero for the imitator) the unsuccessful firm imitates the inventor and the consumer surplus in sector s is positive.




Figure 2.1: Action profiles (k_1, k_2) for different combinations of beliefs. In the dark region the socially optimal action profile $(1, 1)$ differs from the strategic one $(0, 1)$ or $(0, 0)$ for parameter values $\delta = 0.8$, $c = 0.3$, $\lambda_1 = 0.2$, $\lambda_2 = 0.25$, $\underline{x} = 0.75$ and $\tilde{x} = 0.5$.

Strategic R&D investment

Subsequently we characterize the symmetric equilibrium when both firms are in the same sector so that intersectoral spillovers are irrelevant. In this equilibrium both firms choose $k_{s,t} = 1$ above a certain threshold belief \bar{p}_s and $k_{s,t} = 0$ for beliefs below $\underline{p}_s > p_s^a$. In between each firm invests $k_{s,t} \in (0, 1)$ to make the other firm indifferent towards its level of R&D investment.

Proposition 2.1. *If both firms are in the same sector $s = 1, 2$ in the symmetric equilibrium each firm chooses*

$$k_{s,t}^* = \begin{cases} 1 & \text{for } p_{s,t} \in [\bar{p}_s, 1], \\ \frac{-(1-\delta)c + \delta p_{s,t} \lambda_s \underline{x}_s \alpha}{\delta p_{s,t} \lambda_s^2 \frac{1+\gamma_s}{2} \underline{x}_s \alpha} & \text{for } p_{s,t} \in (\underline{p}_s, \bar{p}_s), \\ 0 & \text{otherwise,} \end{cases}$$

for $t \leq T$, where $\underline{x}_1 = 1$, $\underline{x}_2 = \underline{x}$,

$$\underline{p}_s = \frac{(1-\delta)c}{\delta \lambda_s \underline{x}_s \alpha}, \text{ and } \bar{p}_s = \frac{(1-\delta)c}{\delta \lambda_s (1 - \lambda_s \frac{1+\gamma_s}{2}) \underline{x}_s \alpha}.$$




Figure 2.2: R&D investment in equilibrium if both firms are in sector 1 for $\delta = 0.8$, $c = 0.3$, $\alpha = 0.95$, $\lambda_1 = 0.2$ and $\gamma_1 = 1$.

Two failed experiments from \bar{p}_s yield a posterior belief below \underline{p}_s . Similarly, if both firms choose $k_{s,t}^* \in (0, 1)$ for beliefs $p_{s,t} \in (\underline{p}_s, \bar{p}_s)$, the posterior belief is below the lower cutoff belief in the absence of a success. This implies that in the symmetric equilibrium there is at most one time period in which firms choose $k_{s,t}^* \in (0, 1)$. Further, as future payoffs are discounted firms do not have an incentive to postpone experimentation. Figure 2.2 depicts the relationship between the belief and R&D investment in equilibrium.

R&D investment increases if firms become more patient or if the profit from a discovery rises. An increase in the costs of R&D increases the cutoff and hence decreases investment. While the lower cutoff belief is decreasing in the success probability λ_s , the effect on the upper cutoff is ambiguous. The upper cutoff decreases in λ_s if intrasectoral spillovers are small. Moreover, we see that \bar{p}_s is increasing in γ_s , which implies that R&D investment is lower the higher γ_s . As it becomes easier to imitate the success of the other firm, each firm reduces its own investment in equilibrium. Compared to the monopolist firms invest less, i.e., $\underline{p}_s > p^a$, because competition in the product market reduces the expected profits from a discovery (note that $\underline{p}_s = p^a$ for $\alpha = 1$).

Welfare-maximizing R&D investment

To compare the R&D investment of strategic firms to the socially optimal investment, we also restrict the social planner to symmetric strategies. If both firms are in sector s , welfare is given by $W_s = \psi_s + \underline{x}_s \alpha(1 + \gamma_s)$, where $\underline{x}_s \alpha(1 + \gamma_s)$ is the producer surplus if each firm sells one unit.

Proposition 2.2. *The R&D investment of a welfare-maximizing social planner, who is restricted to symmetric strategies, is given by*

$$k_{s,t}^w = \begin{cases} 1 & \text{for } p_{s,t} \in [\bar{p}_s^w, 1], \\ \frac{-(1-\delta)c}{\delta p_s W_s \lambda_s^2} + \frac{1}{\lambda_s} & \text{for } p_{s,t} \in (\underline{p}_s^w, \bar{p}_s^w), \\ 0 & \text{otherwise,} \end{cases}$$

for $t \leq T$, where

$$\underline{p}_s^w = \frac{(1-\delta)c}{\delta \lambda_s W_s}, \text{ and } \bar{p}_s^w = \frac{(1-\delta)c}{\delta \lambda_s (1-\lambda_s) W_s},$$

and $W_s = \underline{x}(1 + \alpha)$ for $\psi_s = \underline{x}(1 - \alpha\gamma_s)$ and $W_s = \underline{x}\alpha(1 + \gamma_s)$ for $\psi_s = 0$.

Comparing the lower cutoff belief of a social planner to the respective threshold used by strategic firms, we see that the social planner experiments until a more pessimistic belief after a history of failures. This is independent of the specification of consumer surplus. Furthermore, socially optimal R&D investment is (weakly) increasing in intrasectoral spillovers γ_s . This is the main difference to strategic investment for which we obtained the opposite result. While the social planner invests more if both firms can benefit from one successful discovery, for strategic firms the incentive to free-ride increases and investment decreases.¹⁰

The upper cutoff belief of the social planner lies below the one for strategic firms if intrasectoral spillovers are high or the success probability λ_s is small. For high levels of γ_s , for example $\gamma_s = 1$, $k_s^w \geq k_s^*$ for all $p_s \in [0, 1]$. In this case the probability of a discovery is higher under the socially optimal R&D investment than if firms choose their R&D investment strategically. In a winner-takes-all competition ($\gamma_s = 0$), welfare-maximizing investment can be higher or lower than strategic investment, depending on the parameters of the model.

¹⁰See e.g. Bolton and Harris (1999), Heidhues et al. (2015), or Keller et al. (2005) for more details on the free-riding problem in games of strategic experimentation.

Let us briefly summarize the main points of this section. If firms are in different sectors intrasectoral spillovers are irrelevant, but intersectoral spillovers affect R&D investment. The firm in the independent sector does not invest enough from the perspective of a social planner, as it ignores the effect of its investment on profits in the dependent sector. When firms are in the same sector, the complementarity is irrelevant and intrasectoral spillovers determine the difference between strategic and welfare-maximizing investment. For high intrasectoral spillovers welfare-maximizing R&D investment exceeds strategic investment as strategic firms free-ride on the R&D investment of the competitor.

2.4 Selecting a sector

Now we want to analyze the effects of inter- and intrasectoral R&D spillovers on the selection of research lines, that is, when firms choose one of the two sectors in the beginning of the game. Firms are not allowed to change their decision later on. Together with the assumption that firms cannot invest simultaneously in both research lines, this reflects the idea that selecting a certain research line is associated with high sunk costs, for instance, because specialized equipment has to be bought or staff has to be trained. R&D investment after sector choice is assumed to be given by the optimal strategic R&D investment as described in Lemma 2.1, Lemma 2.2 and Proposition 2.1.¹¹

First, consider the choice of research lines in $t = 0$ for strategic firms. The expected profit of firm i if both firms are in sector s is

$$\sum_{t=1}^T \delta^{t-1} k_{s,t}^* \left(p_s \left(\delta \alpha (1 + \gamma_s) \lambda_s \left(1 - \frac{k_{s,t}^* \lambda_s}{2} \right) - (1 - \delta)c \right) \prod_{r=1}^{t-1} (1 - k_{s,r}^* \lambda_s)^2 - (1 - \delta)c(1 - p_s) \right).$$

The first term describes the expected profit if the state of the world is good, while the last term represents the bad state of the world. If the state of the world is bad, no discovery will occur and the firm pays the costs of R&D until it stops to invest. If firms invest in different sectors the firm in the independent sector obtains

$$\sum_{t=1}^T \delta^{t-1} k_{1,t}^a \left(p_1 (\delta \lambda_1 - (1 - \delta)c) \prod_{r=1}^{t-1} (1 - k_{1,r}^a \lambda_1) - (1 - \delta)c(1 - p_1) \right),$$

¹¹The main results stated in Propositions 2.3 and 2.4 do not change if R&D investment after sector choice is the socially optimal R&D investment. This only shifts the intervals of beliefs in which the social planner decides differently from the firms.

while the firm in the dependent sector receives

$$\sum_{t=1}^T \delta^{t-1} k_{2,t}^a \left(p_2 (\delta \lambda_2 (\underline{x} + \tilde{x} p_1 q_t(p_1)) - (1 - \delta)c) \prod_{r=1}^{t-1} (1 - k_{2,r} \lambda_2) - (1 - \delta)c(1 - p_2) \right).$$

Firms are more likely to choose the same sector s , the higher p_s and x_s and the higher α and γ_s . The belief and the expected profit x_s increase the likelihood that a firm invests in a given sector independently of strategic considerations. The other two parameters, α and γ_s , reflect how competition affects the firms' decisions. If competition is fierce and imitation difficult, i.e., if α and γ_s are both low, agents are more likely to invest in different sectors. Intersectoral R&D spillovers increase the likelihood that a firm invests in the dependent sector instead of joining a firm in the independent sector.

Let us now compare which research line is chosen by strategic firms to the choice of a social planner. In the sequel we assume that the social planner can only influence the selection of a research line, but not R&D investment after sector choice. Let $u_{s,n}$ denote the individual expected payoff of a firm in sector s if the other firm is in sector n and $W_{s,n}(p_s, p_n)$ denotes welfare if one firm is in sector s and the other in sector n for $s = 1, 2$ and $n = 1, 2$. First, we consider the case where consumer surplus is decreasing in γ_s and welfare does not depend on γ_s , i.e., $\psi_s = \underline{x}_s(1 - \alpha\gamma_s)$ and $W_s = \underline{x}_s(1 + \alpha)$. The results for this case are summarized in Proposition 2.3. Second, we analyze the case where consumer surplus $\psi_s = 0$ (see Proposition 2.4).

If welfare does not depend on γ_s , the social planner prefers joint research in the independent sector for beliefs, where strategic firms research in different sectors. Whether there exists an interval of beliefs for which the social planner prefers joint research in the dependent sector depends on the level of intersectoral spillovers. If intersectoral spillovers are high, i.e., if \tilde{x} is large, the social planner lets the firms pursue different lines of research for beliefs at which the strategic firms both invest in the dependent line.

Proposition 2.3. *Suppose one firm chooses the line of research in $t = 0$ for a given choice of the other firm and consumer surplus is given by $\psi_s = \underline{x}_s(1 - \alpha\gamma_s)$. Then the strategic choice coincides with the choice of a social planner except that for all $\gamma_s \in [0, 1]$*

- (i) *for sufficiently optimistic priors, $p_1 \geq \bar{p}_1$, the social planner lets the firm join the independent sector for beliefs where a strategic firm starts research in the dependent sector;*

(ii) if \tilde{x} is small the social planner lets the firm join the dependent sector for beliefs at which a strategic firm starts research in the independent sector, and if \tilde{x} is large the social planner lets the firm start research in the independent sector for beliefs at which a strategic firm joins the dependent sector.

Part (i) of Proposition 2.3 states that there exists an interval of beliefs for which the social planner prefers if the firm joins the independent sector while a strategic firm starts research in the dependent sector. The reason for this is the following: If both firms invest in R&D in the same sector, they benefit from each others' investment as this increases the probability of a discovery. A firm choosing the sector strategically does not take the positive effect of its R&D investment on the other firm and consumers into account. However, when deciding whether to start research in the dependent sector both the strategic firm as well as the social planner account for the positive impact of R&D investment in the independent sector on expected profits in the dependent sector. Hence, intersectoral spillovers do not lead to any difference between the choice of the strategic firm and the social planner in a situation where one firm is already located in the independent sector and the other firm decides whether to join. Differences in the selection of research lines are then driven by intrasectoral spillovers and consumer surplus. Note that for pessimistic priors p_1 close to p_1^a each firm invests very little into R&D and joint research does not increase the probability of a success. In fact the probability of a discovery if both firms invest little in R&D is lower than if one firm invests its entire resource. Thus, for pessimistic beliefs the social planner prefers diversification over joint research in the same sector.

In part (ii) the firm decides whether to join the dependent sector. It depends on the complementarity between sectors \tilde{x} whether the social planner is more likely to start research in the independent sector. If the benefit associated with two complementary successes is high, there exists an interval of beliefs for which the social planner lets the firm start research in the independent sector while a strategic firm joins the dependent sector, and vice versa if this benefit is rather small. This means that the social planner faces a trade-off between letting firms experiment jointly in the dependent sector, which increases the probability of a discovery, and experimenting in different sectors to potentially exploit the complementarity.

If welfare is increasing in the level of intrasectoral spillovers, it depends on the combination of inter- and intrasectoral R&D spillovers whether the social planner is more likely to let firms research jointly in the same sector.

Proposition 2.4. *Suppose one firm chooses the line of research in $t = 0$ for a given choice of the other firm and consumer surplus is given by $\psi_s = 0$. Then the strategic choice coincides with the choice of a social planner except that*

- (i) *for sufficiently optimistic priors, $p_1 \geq \bar{p}_1$, there exists an interval of beliefs for which the social planner lets firms research jointly in the independent sector while a strategic firm starts research in the dependent sector if $\gamma_1 = 1$ and vice versa if $\gamma_1 = 0$;*
- (ii) *a strategic firm joins the dependent sector for beliefs where the social planner lets firms research in different sectors if $\gamma_2 = 0$. For $\gamma_2 = 1$ the social planner prefers joint research in the dependent sector for beliefs where the strategic firm starts research in the independent sector if and only if \tilde{x} is small.*

If consumer surplus is zero, the decision of the social planner and strategic firms differs because of the diverse influence of intrasectoral spillovers. For $\gamma_1 = \gamma_2 = 0$ the social planner lets firms research in different sectors for beliefs where strategic firms enter the same sector. If only one firm can obtain the profits from a discovery there exists an interval of beliefs where aggregate expected profits are higher if firms do not work on the same research line and thereby compete for the breakthrough. In case of perfect spillovers, $\gamma_s = 1$, the social planner prefers joint research in the independent sector for beliefs where strategic firms pursue different lines of research. Joint research in the dependent sector is, however, only preferred if \tilde{x} is small. In this case specialization in which all research efforts are concentrated in one field yields higher expected payoffs than investing in two complementary sectors. If the extra payoff is small firms are better off increasing the probability of success in the dependent sector than exploiting the complementarity.

2.5 Intersectoral R&D spillovers and intellectual property rights

Finally, we want to explore how firms can be induced to select the socially optimal research line. We assume that the social planner can design IPR, which we interpret as policies that determine the level of intrasectoral spillovers γ_s . As a measure of IPR γ_s can, for instance, be influenced through changes in the duration or strength of patents, or the transparency of the patent system. While a rise in patent duration or strength corresponds to a decrease in γ_s , an increase in transparency reduces the likelihood of duplication, which can be modeled through an increase in γ_s .

We are interested in finding γ_s such that the decision of strategic firms which research line to pursue is aligned with the decision of a social planner. As the exact value of this γ_s depends on how long firms experiment in equilibrium, we focus on the difference between IPR designed for two independent sectors and IPR that take the complementarity into account. IPR that guarantee the socially optimal selection of research lines in $t = 0$ for complementary sectors are denoted by $\tilde{\gamma}_s^*$ and for independent sectors by γ_s^* .

Proposition 2.5. *Intellectual property rights that ensure that in $t = 0$ strategic firms select the same research line as a social planner are such that $\gamma_1^* = \tilde{\gamma}_1^*$ and*

- (i) $\gamma_2^* > \tilde{\gamma}_2^*$ if consumer surplus is decreasing in γ_2 while welfare does not depend on γ_2 ;
- (ii) $\gamma_2^* < \tilde{\gamma}_2^*$ if consumer surplus does not depend on γ_2 , while welfare is increasing in γ_2 .

Intrasectoral spillovers in the independent sector that align the socially optimal choice of sectors with the strategic choice are unaffected by the complementarity. This is intuitive as a firm deciding between joining the independent sector or starting research in the dependent sector benefits from the complementarity and hence takes the positive impact of investment in the independent sector on the dependent sector into account. For sector 2 it depends on the impact of intrasectoral spillovers (or IPR) on consumer surplus and welfare, whether they are higher or lower for dependent sectors. Due to the intersectoral spillovers research in different sectors becomes more attractive compared to joint research in the dependent sector. If consumer surplus is decreasing in γ_2 , the social planner prefers

research in different sectors and IPR are designed such that it is more difficult to imitate an invention in the dependent sector if there is a complementary. Hence, IPR that induce firms to select the socially optimal line of research encourage firms to work in different sectors. If welfare is increasing in γ_2 , imitating an invention is facilitated to encourage firms to both research in the dependent sector.

2.6 Discussion

In this chapter we analyzed a model of innovation in which the profitability of an invention depends on the success of another research project. First, we showed that the number of time periods until firms reach their stopping belief determines whether the complementarity encourages additional R&D investment. Second, we compared strategic and welfare-maximizing R&D investment for a given choice of sectors. In addition to the free-riding inefficiency that occurs for high intrasectoral spillovers, there is another source of inefficiency. If firms work on different research lines, the firm working in the independent sector ignores the positive effect of its investment on profits in the dependent sector.

After that we analyzed the impact of inter- and intrasectoral R&D spillovers on the decision of a firm which research lines to pursue. The social planner prefers more joint research in the independent sector than strategic firms, but more joint research in the dependent sector is only preferred if intersectoral spillovers are small. There exists a trade-off between working in different sectors to exploit the complementary and experimenting jointly which increases the probability of success.

Further, we investigated the possibility to encourage firms to select the socially optimal line of research in the beginning of the game. Intrasectoral R&D spillovers in the dependent sector that align the socially optimal and the strategic choice are lower in case sectors are dependent if consumer surplus is decreasing in the level of intrasectoral spillovers, that is, it is more difficult to imitate an invention in the dependent sector. The complementary leads to an interval of beliefs in which it is socially optimal to work in different sectors and exploit the complementary, but strategic firms enter the same sector. By decreasing intrasectoral spillovers in the dependent sector joint research becomes less

attractive and the strategic firm is more likely to start research in the independent sector. If welfare is increasing in the level of intrasectoral spillovers, imitation in the dependent sector is facilitated to foster joint investment.

Our results emphasize that research and innovation are influenced by both inter- and intrasectoral spillovers, and especially their combination determines which research lines are and should be pursued. Based on the report of the FTC (2011), the U.S. government might consider implementing policies to increase the transparency of patents in the IT sector: on the one hand to encourage investment within the sector, and on the other hand to create positive spillovers to other sectors. Our model suggests that investment in the IT sector can additionally be fostered by decreasing intrasectoral spillovers in a dependent sector (e.g., medicine) to discourage firms from jointly pursuing research in the dependent sector and encourage them to research in different fields. Hence, strong patent protection, long patent periods or financial incentives solely available to inventors in the dependent sector might lead to the same desired effects as an increase in transparency in the independent sector. Such measures at the same time also increase strategic R&D investment within the dependent sector and hence the likelihood of a discovery.

In this chapter we focused on positive intersectoral R&D spillovers in which profits increase due to a complementary discovery. The opposite case, however, can arise as well. That is, new inventions in one sector might have a negative impact on the profitability of another invention. Negative intersectoral spillovers can be modeled by reversing the order of \bar{x} and \underline{x} . This implies that \tilde{x} is negative and the results in Section 2.3 are reversed. More precisely, if a firm in the independent sector makes a discovery, the firm in the dependent sector experiments less. Similarly, under the socially optimal R&D investment firm 1 experiments less when it takes the negative impact of its investment on firm 2 into account. Also the results of Section 2.4 are reversed. Negative intersectoral spillovers strengthen the result that the social planner prefers firms to experiment jointly for beliefs where strategic firms separate. Further, the social planner faces a trade-off between letting firms compete for a discovery in the same sector and separating, where the firm in the dependent sector suffers in case both make a discovery. As for positive spillovers it depends on the magnitude of the profit reduction and success probabilities which effect dominates.

Chapter 3

The transparency of the patent system and its impact on innovation

3.1 Introduction

Innovation benefits society by creating novel products and processes that raise the standard of living, meet unsatisfied needs, and offer solutions to society's challenges in areas such as energy, health or economic growth. The main goal of the patent system is to foster innovation by granting exclusivity rights to the inventor, taking into account that innovation is a complex process which can be expensive, risky and highly unpredictable (The Federal Trade Commission [FTC], 2011). To promote innovation and enhance consumer welfare the patent system and competition policies have to be synchronized. One area of patent law that affects how well the patent system and competition policy work together is notice, i.e., *"how well a patent informs the public of what technology is protected"* (FTC, 2011, p. 2).

According to the FTC (2011) a clear patent notice promotes innovation by spurring collaboration and by helping firms to identify relevant technologies. A patent notice that lacks clarity might be unable to fulfill these tasks. In particular, uncertainties regarding the scope or content of a patent might cause firms to hesitate to invest in innovative activities or encourage them to engage in expensive clearance searches. Besides this, the risk for post-launch patent assertions and litigation is high and, moreover, a clear patent notice reduces wasteful duplication of research efforts. For these reasons the FTC (2011)

identifies the need to improve the transparency of the patent notice.

While there have been numerous discussions in the economics literature whether, when and what kind of patents are necessary to appropriate R&D investment and promote innovation, limited attention has been paid to issues arising from uncertainty associated with granted patents. This is somewhat surprising given that the incentives to invest in R&D are determined by the reliability of patent protection. Thus, to understand the incentives to invest in risky activities it is necessary to investigate the underlying mechanisms explaining how R&D investment reacts to uncertainties in the patent system. To this end, we develop a theoretical model to analyze how the level of transparency affects R&D investment and to find out under which conditions more transparency indeed increases welfare. Once we know how strategic firms react to changes in the level of transparency, we can analyze whether the patent office can increase welfare by changing the level of transparency and characterize the optimal information disclosure policy of the patent office.

To answer these questions we consider a dynamic model of R&D investment in which two firms invest in R&D of uncertain quality (i.e., it is not clear whether research can lead to a discovery and if so, when a discovery will occur). The firms choose a stopping time at which they irrevocably stop investing if they did not observe a success up to this point in time. Each firm's investment decision is private, i.e., unobservable by the competitor. Upon a discovery a firm files a patent application at the patent office. In a transparent system the patent office grants the patent and releases a notice that informs the public (that is, the unsuccessful firm) about the discovery and the patent. In an intransparent system the content of a patent can lack clarity (implying that firms might not be able to identify patents relevant to them) and the scope of protection can be uncertain (meaning that it is not clear whether new inventions lead to a patent and whether new inventions infringe the existing patent).

First, we characterize the optimal stopping times of a monopolist and of firms that cooperate in R&D. In both cases only intransparencies regarding the scope of a patent affect the firms' investment decisions, as in the absence of competition firms always know which patents exist. A lack of clarity regarding patent scope reduces the expected profit

of a patent and consequently R&D investment and welfare are increasing in the level of transparency. When firms compete for patents, different degrees of R&D spillovers are possible. First, we consider a winner-takes-all competition in which only the first discovery obtains a patent and the game ends afterward. If no follow-up inventions are possible, intransparencies regarding the scope of the patent are irrelevant and the stronger firm (i.e., the firm that is more likely to make a discovery) invests more in R&D. Furthermore, the R&D investment of the weaker firm is increasing in the level of transparency, while for the stronger firm it depends on the difference in the firms' R&D productivities. If the two firms are relatively equally strong, the R&D investment of both firms increases in the level of transparency. If, however, the stronger firm is considerably more likely to innovate, her R&D investment is decreasing in the level of transparency. The reason for this is the following. The probability that the weaker firm has an unobserved patent is higher the less transparent the system. At the same time a lack of transparency decreases the R&D investment of the weaker firm and thereby the probability that this firm has an unobserved patent. Depending on which of these two effects dominates, the R&D investment of the stronger firm is increasing or decreasing in the level of transparency.

Second, we consider positive spillovers meaning that the unsuccessful firm can benefit from an invention of the innovator. In this case there does not necessarily exist a (unique) pure strategy equilibrium for certain levels of transparency. While for low and high levels of transparency there exists a unique equilibrium, for intermediate levels no equilibrium in pure strategies exists or two equilibria exist. More precisely, for certain levels of transparency each firm wants to invest in R&D if the competitor does not invest. As firms cannot observe each others' R&D investment, it is possible that both invest, none of them invests or each one of them. In the unique equilibrium the firms' stopping times are non-monotonic in the level of transparency and maximal in a perfectly intransparent patent system in which firms never observe a patent by the rival.

The findings of the FTC further suggest that firms try to influence the clarity of the patent notice through their patent application. Hence, we are interested in knowing what happens if firms can influence the level of transparency, and whether firms have an incentive to do so, e.g., by obscuring their patent applications. In a winner-takes-all competition the weaker firm has no incentive to change the level of transparency, while the

stronger firm wants to be as intransparent as possible to discourage the R&D investment of the weaker firm. If firms can profit from an invention of the competitor, both firms want their competitor to invest as long as possible and select the level of transparency so as to encourage the other firm to invest. As each firm's investment is maximal under full intransparency, high levels of intransparency prevail.

Furthermore, we want to know when the patent office (acting as a benevolent social planner) prefers full transparency. In a winner-takes-all competition the patent office prefers full transparency if the firms are similar in terms of their R&D productivities (which implies that the R&D investment of both firms is increasing in the level of transparency). However, if one firm is considerably more likely to make a discovery and additionally also invests more in R&D in equilibrium, full transparency is not necessarily optimal. A lack of transparency increases welfare if the R&D investment of the stronger firm is decreasing in the level of transparency, while the weaker firm's R&D investment increases. In general, the optimal level of transparency depends on the difference between the expected costs of R&D (which are decreasing in the level of transparency) and the value of a discovery. The higher the value of a discovery, the more the patent office wants to encourage the firms to invest in R&D and thereby increase the chances of an innovation. For perfect positive spillovers, the firms' R&D investment is maximal in a perfectly intransparent system. Hence, the patent office prefers high levels of intransparency as long as the consumer surplus is sufficiently high compared to the costs of R&D.

Finally, we consider the case of sequential innovation, where the first discovery paves the way for follow-up inventions (as e.g., improvements of the original technology). This means that intransparencies regarding patent scope and patent content affect R&D investment simultaneously. Depending on the value of the follow-up invention, sequential innovation is similar to positive spillovers (if this value is high), or to a winner-takes-all competition (if this value is low). If firms are uncertain whether new discoveries will lead to a patent because the patent office might see them as a duplication of the existing technology, the likelihood that firms invest in R&D as well as the expected value of the follow-up invention decrease. The risk of ex post legal disputes on the other hand increases the likelihood that the inventor invests in R&D to obtain a second discovery

despite its negative effect on the expected profit of the firm. By investing in R&D the inventor might be able to prevent the imitator from infringing the first patent. A similar observation can be made for a lack of clarity regarding patent content. In the presence of positive R&D spillovers firms are more likely to invest in R&D for follow-up inventions if the probability of not observing a patent of the competitor is high, because firms cannot as easily free-ride on the R&D investment of the competitor.

A thorough economic analysis of R&D investment has to take into account the risky nature of innovation as well as uncertainties stemming directly from the patent system. Intellectual property rights as an instrument to promote innovation heavily depend on how reliable those instruments are and ignoring possible uncertainties may lead to misleading conclusions. For example, if the inventor has to fear infringement, the starting situation for follow-up inventions for the inventor and the imitator is not the same and their incentives to invest or obscure their patent application differ. This in turn implies that R&D investment for the original discovery is affected and contrary to what one might conjecture, R&D investment and welfare are not inevitably higher if this risk is eliminated. Our results suggest that certain types of uncertainties can actually encourage R&D investment. Thus, in any attempt to increase the transparency of the patent system, it is important to carefully consider different types of intransparencies separately and how each of them interacts with the market situation (e.g., the heterogeneity of the firms, competition and R&D spillovers).

3.1.1 Related literature

In a broader context, this chapter is related to the literature studying the impact of intellectual property rights on innovation.¹ In particular, it relates to work focusing on optimal disclosure of information, i.e., whether or when firms patent and thereby disclose novel technologies. For example, Scotchmer and Green (1990) study how the stringency of the novelty requirement affects the pace of innovation, while more recently, Hopenhayn and Squintani (2016) analyze optimal patents with respect to the timing of information disclosure. So far, optimal information disclosure and the timing of information disclosure

¹See, e.g., Rocket (2010) for a survey.

were studied under the presumption that information disclosure and patent protection are perfect in the sense that there is no uncertainty associated with the patent once it is granted. Only limited attention has been dedicated to questions related to the transparency of the patent system. One exception is Lemley and Shapiro (2005), who discuss uncertainty associated with intellectual property rights, propose reforms to reduce this uncertainty and finally discuss the effects of uncertainty on the incentives of firms to settle disputes or litigate. Bessen and Meurer (2008) discuss problems of uncertainty (or intransparency) of intellectual property rights as well as reform suggestions to improve patent notice. Similar to the FTC (2011) the authors argue that clarity is central to attain efficiency. However, none of these papers investigates the linkages between (strategic) R&D investment, uncertain patent protection and competition in detail.

Furthermore, this chapter is related to the growing literature on strategic experimentation, which offers a suitable framework for modeling situations where uncertainty plays a crucial role in the innovative process.² The modeling framework employed in this chapter is closely related to the model of Bonatti and Hörner (2011), where innovations arrive in the form of fully revealing breakthroughs. In Bonatti and Hörner (2011) discoveries are public while R&D investment (or effort) is private.³ Many models of strategic experimentation (including Bolton and Harris, 1999, Keller et al., 2005, Keller and Rady, 2010, Bonatti and Hörner, 2011, Bimpikis and Drakopolous, 2014, Heidhues et al., 2015) study a team problem in which the discovery of one team member has perfect positive spillovers to the other members of the team. In these models severe forms of free-riding arise due to the public goods character of information. More recently, competition has been introduced into these models to study situations in which the inventor has some advantage compared to the unsuccessful imitators. Examples include Besanko and Wu (2010, 2013), Acemoglu et al. (2011), Das (2013), Akcigit and Liu (2015) and Wong (2015), where the last three also allow for heterogeneity across agents. Akcigit and Liu (2015) study private R&D investment of firms that compete for patents when an arrival can either lead to a patent (which is publicly observable) or to a dead-end (which is only privately ob-

²See, e.g., Bergemann and Välimäki (2008) for a survey.

³In Bonatti and Hörner (2015), R&D investment is also private, but information arrives in the form of fully revealing publicly observable breakdowns.

served). Firms do not only compete on a risky research line, but also on a safe line and thus have an incentive to keep any dead-end findings secret to avoid competition on the safe line. Wong (2015) studies optimal patent protection in a setting without informational spillovers and shows that strict patent protection can lead to duplication of R&D, while loose patent protection fosters free-riding.

Subsequently, we study a model in which two heterogeneous firms compete in R&D. Additional to R&D investment being private as in Bonatti and Hörner (2011), also discoveries may be private in case the patent system lacks transparency. In contrast to most papers on strategic experimentation future payoffs are not discounted.

3.2 Model

3.2.1 Investment in uncertain research projects

There are two firms $n = i, j$ that can invest in R&D in continuous time $t \in [0, \infty)$. R&D investment is risky, meaning that the research in which firms can invest is of one of two types, $\theta = \{0, 1\}$. A bad research project, $\theta = 0$, never leads to a discovery. In a good project, $\theta = 1$, discoveries occur at exponentially distributed random times, where the arrival is independent across firms. Investment in research is costly and the marginal costs of R&D are denoted by c , so that firm n pays $c dt$ if she invests in R&D in the time interval $[t, t + dt)$. The probability that a firm makes a discovery is $\lambda_n \theta dt$ if the firm invests in R&D, where $\lambda_n > 0$ is a constant that can be interpreted as firm-specific R&D productivity and is known to both firms. In what follows we assume that $\lambda_i > \lambda_j$ and hence refer to firm i also as the stronger firm and to firm j as the weaker firm. An arrival represents a discovery or invention and the reward of a discovery (that is, the profit a firm obtains) is determined by the value of the patent.

3.2.2 The value of a patent and R&D spillovers

After a discovery, the firm files a patent application at the patent office.⁴ Once the patent office grants patent $y = 1, 2, \dots$ the firm obtains a flow profit π_y for length τ , so that the value of patent y at the time of approval is $\pi_y\tau$. Furthermore, the discovery of one firm might have positive R&D spillovers to the other firm that determine whether an imitator can profit from the discovery as well. The unsuccessful firm obtains a flow profit of $\omega_y \geq 0$ also for length τ ⁵ and we distinguish two possibilities. First, $\pi_y = \omega_y$, which means we have perfect positive R&D spillovers, where a discovery by firm n can be profitably used by firm $-n$ as well. For example, firms might sell different products using similar technologies and consequently are able to serve different markets. Second, $\omega_y = 0$ meaning that firms face a winner-takes-all competition, where only the inventor profits from a discovery. For example, a discovery might lead to the formation of a natural monopoly. The total profit of the first discovery in a transparent patent system is given by

$$W_1 = \pi_1\tau + \omega_1\tau. \quad (3.1)$$

After the first patent is granted, firms can continue to invest in R&D to obtain further patents. Subsequent innovations occur on the same research line, that is, innovations are sequential and one discovery builds on the previous discovery.⁶ We make the simplifying assumption that there are at most two patents granted for discoveries in one research area so that the game ends after the patent office grants the second patent.

3.2.3 Beliefs and informational spillovers

Firms have a belief about the research line being viable and both firms start with a common prior p . An arrival on the research line reveals that it is good and the firm's belief

⁴For simplicity we assume that without applying for a patent, the new technology is immediately imitated by many firms and payoffs are driven to zero. Hence, upon discovery firms always apply for a patent. Moreover, an application is filed at the time of discovery and the patent office decides immediately whether the patent is granted.

⁵This assumption is w.l.o.g. as we assume that the imitator obtains this payoff independently of the level of transparency and the results would not change if we assumed $\tau_\pi \neq \tau_\omega$ instead.

⁶New inventions on the same line might replace existing technologies. As this happens independently of the level of transparency of the patent system, we abstract from this possibility here and assume that there are no negative externalities from a new discovery on existing technologies.

jumps to 1. As long as firms invest in research without a success, they become increasingly pessimistic about the feasibility of the project. Firms are said to experiment if they invest in research before knowing the type of the project. To exclude trivial cases in which investing is never optimal we assume that the ex ante expected benefits of R&D exceed the costs.

Assumption 1. $\pi_1 \tau p \lambda_n > c$ for $n = i, j$.

The types of the research project of firm i and firm j are assumed to be perfectly positively correlated. This means that a discovery of firm i reveals to firm j that the project is good and vice versa. The idea is that two firms working in the same industry, e.g., the pharmaceutical industry, both can potentially work on the same research line, as e.g., the discovery of a new drug against a certain disease. If one firm succeeds this reveals to the other firm that this research line is good and, say, a cure of this disease is achievable. Possible interpretations are that firms use a similar approach or that the unsuccessful firm can easily imitate the approach or technology of the inventor. Moreover, we assume that the first success indicates that a second success is possible. That is, the first arrival is the breakthrough or fundamental discovery that the project is good, while the second arrival can be interpreted as an improvement of the existing technology. In the case of pharmaceuticals this could be a reduction of side effects or better understanding of the right dose.

3.2.4 Information structure and timing

In a perfectly transparent patent system a patent (notice) provides a clear indication of what it protects. Hence, duplications as well as ex post legal disputes are ruled out. Subsequently full transparency is interpreted as follows:

- The patent application is public information and consequently so is a discovery;
- the approval of the patent is public information;
- the inventor knows that her invention is protected and she obtains $\pi_y \tau$ with certainty, that is, there is no risk of ex post legal disputes;

- an imitator can modify her own invention such that a new patent can be obtained. As the imitator is perfectly informed about the extent of protection, she can assess which changes to the original discovery are necessary to avoid a violation of the intellectual property of the inventor. This means that in fact every arrival (discovery) obtains a patent.

Based on the FTC (2011) report two types of intransparencies can arise:

1. *Unclear patent content*: If the content of a patent is unclear, it is difficult for firms to assess what innovations already exist. A transparent and effective patent notice enables firms to identify the patents relevant to them. Even though there are online databases providing this information, especially in the IT sector clearance searches are almost prohibitively expensive, because of the large number of potentially relevant patents. Additionally, the language used is often ambiguous and moreover product life cycles in the IT industry are usually short implying that a thorough search might be too time consuming. The FTC (2011, p. 90) refers to this as "*difficulties in sifting through a multitude of patents*". We capture this intransparency by assuming that a patent is observed by the rival with probability $1 - \varepsilon$. A firm that does not observe a patent of the competitor at the time it is granted only learns about it in case of an own success.
2. *Unclear patent scope*: Even if firms are able to identify all relevant patents, the scope of a patent (or patent claim), which defines the extent of protection, might be unclear. That is, firms exhibit "*difficulties in interpreting the boundaries of issued claims*" (FTC, 2011, p. 81). The main problems reported to the FTC were a lack of clarity of the language, varying nomenclature, functional claiming (i.e., explaining what the invention does rather than what it is), and institutional concerns⁷. We consider two possible consequences of a lack of clarity regarding patent scope. First, a patent holder does not know how well her invention is protected by the patent. That is, with probability α an existing invention loses value if a second innovation obtains a patent, for instance, because the firm has to enforce its intellectual property rights

⁷Litigation as the only available mechanism to test what a patent really covers is highly uncertain, expensive and causes delays.

by law resulting in additional expenses. In this case the profit is reduced to $\pi'_1 < \pi_1$.⁸ Second, a firm with a new discovery does not know whether this invention will obtain a patent. To be more precise, a new patent is granted with probability $1 - \beta$, while with probability β it is seen as a duplication of an existing patent.

Let us now turn to the timing of the game. The game can be divided into two stages:

1. *Stage 1, the experimentation stage:* The firms start with prior belief p and decide how long to experiment. That is, each firm $n = i, j$ chooses a stopping time T_n at which she stops to invest in R&D in case she does not observe a discovery (from her own research or the research of the competitor). Each firm's investment decision is private, i.e., unobservable by the other firm. In a transparent patent system patents are always observed, while in an intransparent system a patent of the competitor is observed with probability $1 - \varepsilon$. This stage ends for a firm if she stops to invest once and for all or if she observes a success. After making a discovery, the firm files an application at the patent office. If the patent is granted, the firm proceeds to the second stage.
2. *Stage 2, follow-up inventions:* The first discovery always obtains a patent. The firm with patent 1 receives the expected profit from this patent and can still research to obtain a second patent which is granted with probability $1 - \beta$. If the unsuccessful firm is aware of the existing patent, she can as well invest in research to obtain the second patent (which is granted with probability $1 - \beta$ and reduces the profit of the inventor with probability α). In case the unsuccessful firm is unaware of the existing patent, she is still in stage 1 (or exited the game) and invests either until reaching her stopping time in stage 1 or until an own discovery. The imitator receives ω_y conditional on knowing about patent y .

The game ends after the second patent is granted, after both firms have reached their stopping times without a discovery in stage 1 or if both firms have unsuccessfully applied for a patent in stage 2.

⁸Note that only the inventor faces the risk of litigation and consequently the imitator obtains ω_y independently of any intransparencies regarding patent scope.

3.2.5 Strategies and equilibrium concept

Firms are risk neutral expected profit maximizers and are restricted to pure stopping strategies. Let $T_n^s \in [0, \infty)$ denote the stopping time of firm n in stage s , meaning that at T_n^s firm n irrevocably stops to invest in R&D conditional on not having observed a success in this stage. A strategy of firm n consists of a pair of stopping times (T_n^1, T_n^2) , one for each stage of the game.⁹ In equilibrium the stopping times jointly maximize the firm's expected profit given the stopping times of the competitor. A firm chooses its stopping time in stage 1 anticipating the optimal decisions in stage 2. The solution concept is subgame perfect Nash equilibrium.

3.2.6 Consumers

So far we only discussed the firm side, ignoring the impact of an invention on consumers. Let ψ_y denote the value of discovery y for consumers. The consumer surplus of a discovery is determined by the difference between its value and its price. We assume that the price equals the profit of the firm (which implies that the marginal costs of production are zero) and the consumer surplus can be written as $CS_y = \psi_y - \pi_y\tau - \omega_y\tau$. The producer surplus for the two firms is the difference between the profit earned from a discovery and the costs of development. As the price (=profit) is merely a transfer between firms and consumers, a social planner who aims to maximize total welfare maximizes

$$E[\psi_y - C_y],$$

where C_y denotes the costs of development.

3.3 The monopolist

Let us start by considering the R&D investment decision of a single firm that acts as a monopolist. In this case there is no competition and consequently intransparencies regarding the content of a patent (i.e., ε) as well as ex post legal disputes (i.e., α) are irrelevant. However, an unclear patent scope makes it difficult for a firm to evaluate

⁹For simplicity of exposition the superscript indicating the stage will be dropped in the remaining chapter and T_n refers to the stopping of firm n in stage 1.

whether additional inventions will lead to a new patent. We start in stage 2 assuming that the monopolist already holds one patent and decides whether to invest in R&D to obtain a second patent.

Stage 2: Suppose firm n holds patent 1. This patent generates a flow profit π_1 for length τ . As described in Section 2 we assume that the type of the first and the second discovery are perfectly positively correlated. The firm knows that the risky project generates positive arrivals ($p = 1$) and compares the expected benefit of the next arrival with the costs of research. As there is no learning or belief updating, investing in R&D in stage 2 for the monopolist yields an expected payoff of

$$v_n^m = -c dt + (1 - \beta)\pi_2\tau\lambda_n dt + (1 - \lambda_n dt)v_n^m.$$

The first term on the r.h.s. are the costs of research the firm pays in the interval $[t, t + dt)$. The second term describes the continuation payoff in case the firm obtains patent 2, which happens with probability $(1 - \beta)\lambda_n dt$. With probability $1 - \lambda_n dt$ the firm does not observe a success and invests again in the next period. The expected payoff from investing until a discovery is given by

$$v_n^m = \frac{\lambda_n(1 - \beta)\pi_2\tau - c}{\lambda_n}. \quad (3.2)$$

If the expected benefit of patent 2 exceeds the costs of R&D, the firm invests in R&D in stage 2. As beliefs do not change, the firm either does not invest at all in stage 2 or invests until a discovery. The expected payoff associated with the second patent is increasing in the arrival rate of the discovery and decreasing in the costs. Besides this, it is decreasing in β , and hence it is less likely that a firm invests the higher the level of intransparency.

Stage 1: The expected payoff of firm n for a given prior p and stopping time T_n is

$$U_n = -(1 - p)cT_n + p(\lambda_n W_n^m - c) \frac{1 - e^{-\lambda_n T_n}}{\lambda_n}, \quad (3.3)$$

where $W_n^m = \pi_1\tau + \max\{0, v_n^m\}$.¹⁰ The first term on the r.h.s. captures the bad state of the world, which occurs with probability $1 - p$. As in this case no discovery can be made, the firm pays the costs of R&D until reaching her stopping time T_n . If the state of the world is

¹⁰In case the expected payoff of R&D in stage 2 is negative ($v_n^m < 0$), firm n does not invest in R&D in stage 2 and obtains only $\pi_1\tau$ after a success in stage 1.

good, the firm pays $-c dt$ and makes a discovery with probability $\lambda_n dt$ in any time interval of length dt until T_n as long as there was no discovery before. Maximizing the expected payoff w.r.t. the stopping time T_n we obtain the first order condition

$$-(1-p)c + p(\lambda_n W_n^m - c)e^{-\lambda_n T_n} = 0. \quad (3.4)$$

At the time of stopping the firm has to be indifferent between the expected payoff of experimenting one last time and the payoff of stopping right now which equals zero. Hence, for a monopolist the optimal time to stop investing in stage 1 equals

$$T_n^m = -\frac{1}{\lambda_n} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_n W_n^m - c} \right) \right]. \quad (3.5)$$

The stopping time T_n^m is increasing in the prior p and W_n^m and decreasing in the costs c . The effect of an increase in the arrival rate λ_n on the stopping time is ambiguous. A higher λ_n leads to a faster updating of beliefs, meaning that the firm becomes more pessimistic in the absence of a success. At the same time a higher λ_n decreases the belief at which a firm stops investing and it is not clear which effect dominates.¹¹

Firm n experiments in stage 1 until it makes a discovery or stops at T_n^m . The stopping time is higher in a transparent system, because the expected benefit of R&D, W_n^m , is decreasing in β . That is, firms invest more and the probability of mistakenly abandoning a good project is smaller the higher the level of transparency. However, this is only true if the firm invests in stage 2 (i.e., if $v_n^m > 0$). If there is no investment in stage 2, the optimal stopping time in stage 1 does not depend on the level of transparency.

In the case of a monopolist who is able to reap the entire consumer surplus, maximizing welfare is equivalent to maximizing expected monopoly profits. Thus, increasing transparency increases welfare. If the monopoly profit lies below the valuation of consumers, increasing transparency increases welfare as R&D investment and thus the probability of an invention are increasing in the level of transparency.

3.4 Research cooperation

Now we turn to the optimal R&D investment of two firms in a research alliance, i.e., firms that cooperatively maximize aggregate expected payoffs. Cooperation is restricted to the

¹¹See also Halac et al. (2013).

research stage meaning that firms remain competitors in the product market. If firms cooperate, arrivals are always revealed as it can never be optimal to hide information about a discovery. Hence, again only intransparencies about patent scope matter and ex post legal disputes are precluded. In contrast to the previous section, positive spillovers might allow the imitator to earn profits from a discovery of the inventor. As before, we start in stage 2, where one firm already holds a patent.

Stage 2: The expected payoff of research in stage 2 in a research alliance where both firms invest is given by

$$v^a = -2cdt + \Lambda(1 - \beta)(\pi_2 + \omega_2)\tau dt + (1 - \Lambda dt)v^a,$$

where $\Lambda = \lambda_i + \lambda_j$ and in contrast to the monopolist profits for the imitator, ω_2 , enter as well. The expected payoff of R&D in stage 2 is

$$v^a = \frac{\Lambda(\pi_2 + \omega_2)\tau(1 - \beta) - 2c}{\Lambda}, \quad (3.6)$$

if both firms invest in R&D, and

$$v_n^a = \frac{\lambda_n(\pi_2 + \omega_2)\tau(1 - \beta) - c}{\lambda_n},$$

if only firm n invests, but the imitator also obtains profits in case of a success. As firm i is stronger in the sense that $\lambda_i > \lambda_j$, it follows that $v_j^a < v^a < v_i^a$. Therefore, in a research alliance only the stronger firm invests in R&D in stage 2 or none of the firms invests in R&D.¹²

Stage 1: Let W^a denote the sum of the total profit in stage 1 and the expected value of research in stage 2 in a research alliance, that is, $W^a = \pi_1\tau + \omega_1\tau + \max\{0, v_i^a\}$. The total expected payoff of firms in a research alliance is given by

$$U_i^a + U_j^a = pW^a(1 - e^{-\lambda_i T_i - \lambda_j T_j}) - (1 - p)c(T_i + T_j) - pc \left(2 \frac{1 - e^{-\Lambda T_{\min}}}{\Lambda} + e^{-\lambda_{\min} T_{\min}} \frac{e^{-\lambda_{\max} T_{\min}} - e^{-\lambda_{\max} T_{\max}}}{\lambda_{\max}} \right), \quad (3.7)$$

where $T_{\min} = \min\{T_i, T_j\}$, $T_{\max} = \max\{T_i, T_j\}$ and λ_{\min} is the arrival rate corresponding to the firm with the earlier stopping time, that is, if $T_j = T_{\min}$, then $\lambda_j = \lambda_{\min}$ and $\lambda_i = \lambda_{\max}$.

¹²Homogenous firms, i.e., when $\lambda_i = \lambda_j$, either both invest or none of the two firms invests in R&D.

The first term is the expected benefit of R&D investment. With the probability that the state of the world is good and there will be at least one discovery, $p(1 - e^{-\lambda_i T_i - \lambda_j T_j})$, the firms obtain W^a . If the state of the world is bad, both firms pay the costs of R&D until reaching their stopping times. Moreover, firms also pay the costs of R&D in the good state of the world as long as no one made a discovery. Here $2\frac{1-e^{-\Lambda T_{\min}}}{\Lambda}$ refers to the time interval in which both firms invest, and the last term describes the time interval in which one of the two firms already stopped investing, while the other firm still continues.

The firms maximize the sum of their expected profits given in (3.7) w.r.t. their stopping times T_i and T_j . As stated in Lemma 3.1 in a research alliance also in stage 1 all R&D investment is carried out by the stronger firm.

Lemma 3.1. *The optimal stopping times of firms in a research alliance are $T_j^a = 0$ and*

$$T_i^a = -\frac{1}{\lambda_i} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_i W^a - c} \right) \right].$$

Similar as in stage 2, only the stronger firm invests in R&D in a research cooperation. Expected payoffs are higher if the stronger firm invests in two consecutive time periods than if both firms invest simultaneously. The rate at which firms update their belief in the absence of a success in a research alliance is the same as for the monopolist firm i . The stopping time T_i^a equals the stopping time of the monopolist firm i , i.e., $T_i^a = T_i^m$ if $W^a = W_i^m$. Note that for $T_i^a = T_i^m$ also the final posterior belief in case there is no discovery is the same in both situation. For $W^a > W_i^m$ cooperating firms invest more in R&D, because the expected benefit of research is higher when both firms can profit from one invention due to the positive spillovers.

In stage 2, if the benefits of R&D exceed the costs only the stronger firm invests and hence similar patterns of R&D investment as for the monopolist are possible. If $v_i^a = v_i^m$, i.e., if the expected payoff a monopolist earns from investing in stage 2 equals the expected payoff that both firms in a research alliance earn, the research alliance and the monopolist firm i are equally likely to invest in R&D. If v_i^a exceeds v_i^m it is more likely that the research alliance invests in R&D and if $v_i^a < v_i^m$, the monopolist is more likely to invest.

If consumer surplus equals zero ($\psi_y = \pi_y \tau + \omega_y \tau$), the research alliance acts socially optimally and as profits are increasing in the level of transparency, welfare is increasing

in the level of transparency. If the consumer surplus is positive, i.e., for $\psi_y > \pi_y\tau + \omega_y\tau$, the more firms invest, the higher is the probability of a discovery and hence more transparency is welfare enhancing.

3.5 Research competition

Commonly firms compete for patents and hence do not (or at least not fully) disclose their research activities publicly. In contrast to R&D cooperation, competing firms choose their R&D investment strategically without observing the R&D investment of the competitor. Furthermore, in an intransparent patent system, a patent is not necessarily revealed to both firms. The unsuccessful firm observes a discovery of the competitor with probability $1 - \varepsilon$, where $\varepsilon \in [0, 1]$. Here $\varepsilon = 0$ implies full transparency where discoveries are revealed publicly immediately. A firm who does not observe the discovery of the competitor only learns about it in case she makes a discovery herself. We assume that a firm is being informed about an existing patent when filing an application at the patent office.

3.5.1 Winner-takes-all competition

Let us start by considering the simplest case in which only one patent can be obtained and there are no positive R&D spillovers between the two firms. That is, the patent race is a winner-takes-all race in which the first discovery receives a patent and the game ends afterwards. If no follow-up inventions are possible, uncertainties about patent scope (α and β) are irrelevant and only a lack of clarity regarding patent content affects R&D investment.

Full transparency

First, we characterize the optimal stopping times for $\varepsilon = 0$. Suppose j follows a stopping strategy with stopping time T_j . We are interested in finding i 's best response. The expected payoff of firm n for a given combination of stopping times is given by

$$U_n = p(\pi\lambda_n - c) \left(\frac{1 - e^{-\Lambda T_{\min}}}{\Lambda} + e^{-\lambda_n T_{\min}} \frac{e^{-\lambda_n T_{\min}} - e^{-\lambda_n T_n}}{\lambda_n} \right) - c(1 - p)T_n, \quad (3.8)$$

where π is the profit associated with the patent, i.e., $\pi = \pi_1\tau$ and $\omega_1 = 0$. At the time of stopping each firm has to be indifferent between experimenting one last time and stopping immediately given the stopping time of the competitor. For firm i this implies

$$-c(1 - p + pe^{-\lambda_i T_i - \lambda_j T_{\min}}) + p\pi\lambda_i e^{-\lambda_i T_i - \lambda_j T_{\min}} = 0. \quad (3.9)$$

The first term represents the costs of R&D the firm pays in case it will indeed experiment in this last instant, which happens with the probability that there was no discovery up to this point in time. The second term is the expected continuation payoff the firm obtains in case of a success. With the probability that the state of the world is good even though no one made a discovery so far, the firm obtains a patent with instantaneous probability λ_i . Similarly, at the time of stopping for firm j the following first order condition has to be satisfied

$$-c(1 - p + pe^{-\lambda_i T_{\min} - \lambda_j T_j}) + p\pi\lambda_j e^{-\lambda_i T_{\min} - \lambda_j T_j} = 0. \quad (3.10)$$

From equations (3.9) and (3.10) the optimal stopping times in equilibrium can be derived. Furthermore, it can be shown that in equilibrium the stronger firm experiments longer, i.e., $T_i > T_j$.

Lemma 3.2. *The optimal stopping times of firms in a winner-takes-all competition with one patent and full transparency ($\varepsilon = 0$) are given by*

$$T_j = -\frac{1}{\Lambda} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_j \pi - c} \right) \right],$$

and

$$\begin{aligned} T_i &= -\frac{1}{\lambda_i} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_i \pi - c} \right) + \lambda_j T_j \right] \\ &= -\frac{1}{\Lambda} \left[\ln \left(\frac{1-p}{p} \right) + \ln(c) - \frac{\Lambda}{\lambda_i} \ln(\lambda_i \pi - c) + \lambda_j \ln(\lambda_j \pi - c) \right]. \end{aligned}$$

In this equilibrium

- (i) the stronger firm invests more in R&D, i.e., $T_i > T_j$;
- (ii) in the absence of a discovery firms learn as much as under cooperation, i.e., the final posterior belief

$$p^f = \frac{pe^{-\lambda_i T_i - \lambda_j T_j}}{1 - p + pe^{-\lambda_i T_i - \lambda_j T_j}},$$

is the same under competition and cooperation as $\lambda_i T_i + \lambda_j T_j = \lambda_i T_i^a + \lambda_j T_j^a$.

Let us compare the optimal stopping times of competing firms to their respective counterparts in the previous sections (assuming a winner-takes-all competition in which only one patent is granted). The weaker firm j invests more than under cooperation but less than as a monopolist, that is, $0 = T_j^a < T_j < T_j^m$. Further, the stopping time of the weaker firm j is independent from the stopping time of the stronger firm i (it does however depend on λ_i), while the optimal stopping time of the stronger firm is decreasing in the stopping time of the weaker firm. The stronger firm i invests less than under cooperation, i.e., $T_i < T_i^a = T_i^m$, where the difference between T_i and T_i^a is determined solely by T_j . As T_i is decreasing in T_j and $T_j > T_j^a$ this implies that $T_i < T_i^a$. Compared to the R&D investment in a research alliance firm j invests too much, while firm i invests too little.

Each firm invests less than as a monopolist, because they can learn from each others lack of success. Not observing a discovery from the opponent reduces the likelihood that the state of the world is good. In equilibrium the firms correctly anticipate each others stopping times. As firm j knows that firm i experiments until T_i , firm j learns more than in autarky, i.e., her final posterior belief is more pessimistic. In particular, both firms have the same final posterior belief at T_i , which equals the posterior of firm i in autarky.¹³

Imperfect transparency

Suppose $\varepsilon > 0$, which means that a firm that did not observe the patent of the competitor at the time it is granted, only learns about it in case of an own discovery. The expected payoff of firm n is given by

$$U_n = p\pi\lambda_n \left(\frac{1 - e^{-\Lambda T_{\min}}}{\Lambda} + \frac{e^{-\Lambda T_{\min}} - e^{-\lambda_n T_n - \lambda_{-n} T_{\min}}}{\lambda_n} \right) - c(1-p)T_n - cp\varepsilon \frac{1 - e^{-\lambda_n T_{\min}}}{\lambda_n} - cp(1-\varepsilon) \frac{1 - e^{-\Lambda T_{\min}}}{\Lambda} - cp(\varepsilon + (1-\varepsilon)e^{-\lambda_{-n} T_{\min}}) \frac{e^{-\lambda_n T_{\min}} - e^{-\lambda_n T_n}}{\lambda_n}. \quad (3.11)$$

¹³Deviations cannot be observed unless one of the two firms makes a discovery after reaching her stopping time T_n . If this happens both firms know that the state of the world is good and the game ends. If firm n stops before reaching her stopping time T_n , the other firm will never know about this deviation and interpret the absence of a discovery mistakenly as a lack of success. Thus, the deviator will be more optimistic than the firm that did not deviate. If a firm continues to invest after reaching her stopping time she will be more pessimistic than the firm that did not deviate.

The first term is the expected benefit of R&D investment, while the remaining terms represent the expected costs.¹⁴ As the costs are increasing in ε , the expected payoff U_n is decreasing in ε , i.e., the lower the level of transparency the lower is the expected payoff for a given combination of stopping times. The optimal stopping times in equilibrium have to satisfy

$$-c(1-p+pe^{-\lambda_i T_i}(\varepsilon+(1-\varepsilon)e^{-\lambda_j T_{\min}}))+p\pi\lambda_i e^{-\lambda_i T_i-\lambda_j T_{\min}}=0, \quad (F^1)$$

and

$$-c(1-p+pe^{-\lambda_j T_j}(\varepsilon+(1-\varepsilon)e^{-\lambda_i T_{\min}}))+p\pi\lambda_j e^{-\lambda_i T_{\min}-\lambda_j T_j}=0. \quad (F^2)$$

Now not observing an invention up to time T_n can either be due to the absence of a success or due to a discovery of the competitor that was unobserved by firm n . Then firm n , being unaware of the competitor's success, continues to invest (pay c), but will never obtain a patent as π can only be realized if there was no discovery.

Solving for the optimal stopping times explicitly as for $\varepsilon = 0$ is not possible, however, in the appendix it is shown that a solution to equations (F¹) and (F²) exists and is unique. In this equilibrium the stronger firm invests more in R&D than the weaker firm. Furthermore, by making use of the implicit function theorem the impact of the level of intransparency ε on the stopping times, that is, $\partial T_i/\partial \varepsilon$ and $\partial T_j/\partial \varepsilon$ can be derived. As stated in Proposition 3.1 the stopping time of the weaker firm is decreasing in ε , while for the stronger firm it depends on the difference in the firms' R&D productivities.

Proposition 3.1. *The optimal stopping times of firms in equilibrium in a winner-takes-all competition with one patent and transparency level ε are uniquely determined by equations (F¹) and (F²). In this equilibrium*

- (i) *the stronger firm invests more, i.e., $T_i > T_j$ for all $\varepsilon \in [0, 1]$;*
- (ii) *the R&D investment of the weaker firm j is decreasing in ε , i.e., $\partial T_j/\partial \varepsilon < 0$;*
- (iii) *the R&D investment of the stronger firm i is decreasing in ε , i.e., $\partial T_i/\partial \varepsilon < 0$ if and only if*

$$(1-e^{-\lambda_j T_j})(e^{-\lambda_i T_i}\Lambda(\lambda_j \pi - c + c\varepsilon) - \varepsilon c \lambda_j) > \lambda_j(\lambda_i \pi - c + \varepsilon c)(1-e^{-\lambda_i T_i})e^{-\lambda_j T_j}. \quad (3.12)$$

¹⁴In the appendix the payoff function is derived and explained in more detail.

The higher ε , the higher is the probability that the opponent has an unobserved patent and hence the lower is the expected profit from experimenting. To be more precise, the expected costs of R&D increase in ε while the expected benefit decreases. If inequality (3.12) is satisfied the R&D investment of both firms is increasing in the level of transparency (or decreasing in the level of intransparency ε). Otherwise the R&D investment of firm i increases in ε . Inequality (3.12) holds if λ_i is close to λ_j and is reversed if λ_i is large relative to λ_j . This implies that if the stronger firm is considerably more likely to make a discovery, her stopping time is increasing in ε . If, however, the two firms are relatively equally strong, then the stopping times of both firms are decreasing in ε . The intuition for this result is the following. The probability that the opponent had an unobserved discovery up to a certain point in time is increasing in ε and this decreases the expected benefit of R&D investment. As a response, firms reduce their investment. However, for firm i at the time of stopping the probability that firm j had an unobserved discovery is higher the more j invested, i.e., the higher T_j . As T_j is decreasing in ε , there are two opposing forces at play. On the one hand the probability that firm j had an unobserved discovery increases in ε . On the other hand, the probability that firm j had an unobserved discovery decreases, because firm j stops earlier the higher ε . If inequality (3.12) is satisfied (meaning if firms are almost equally strong) the former effect dominates the later and vice versa if the difference in the firms' R&D productivities is more pronounced.

Do firms have an incentive to influence the level of transparency?

The FTC (2011) report suggests that firms deliberately try to influence the level of transparency, for example, through vague formulations of their patent application, by describing what the innovation does instead of what it is and so forth. Besides the patent application itself, firms may have other possibilities to disclose more or less information about their research activities (e.g., releasing information about research projects and progress). To find out whether firms have an incentive to influence the level of transparency (e.g., by obscuring the patent application) we have to analyze the impact of the level of intransparency ε on expected payoffs in equilibrium. We assume that firms can solely influence the transparency level (and hence the stopping time) of the competitor. That is, each firm cannot influence overall ε but only the level of transparency faced by

the opponent.¹⁵ Let ε_n denote the level of intransparency faced by firm $-n$ (determined by firm n) and let us see whether firms have an incentive to influence the level of transparency. The expected equilibrium payoffs for a given combination of stopping times T_i and T_j where $T_i > T_j$ are

$$U_i = -c(1-p)T_i + p(\pi\lambda_i - c(1-\varepsilon_j))\frac{1-e^{-\Lambda T_j}}{\Lambda} - cp\varepsilon_j\frac{1-e^{-\lambda_i T_j}}{\lambda_i} + p\lambda_i\pi\frac{e^{-\Lambda T_j} - e^{-\lambda_j T_j - \lambda_i T_i}}{\lambda_i} - cp(\varepsilon_j + (1-\varepsilon_j)e^{-\lambda_j T_j})\frac{e^{-\lambda_i T_j} - e^{-\lambda_i T_i}}{\lambda_i}, \quad (3.13)$$

and

$$U_j = -c(1-p)T_j + p(\pi\lambda_j - c(1-\varepsilon_i))\frac{1-e^{-\Lambda T_j}}{\Lambda} - cp\varepsilon_i\frac{1-e^{-\lambda_j T_j}}{\lambda_j}. \quad (3.14)$$

The weaker firm (i.e., the firm stopping earlier) has no incentive to influence the stopping time of the stronger firm, as T_i has no impact on firm j 's payoffs. To be more precise, the weaker firm does not care about the level of transparency faced by its opponent. The expected profit of the stronger firm i , however, is decreasing in the stopping time of the competitor T_j . Hence, firm i has an incentive to increase the level of intransparency, as T_j is decreasing in ε_i . Firm i prefers to be as intransparent as possible to discourage firm j 's investment.

Proposition 3.2. *Suppose firms face a winner-takes-all competition with one patent only and both firms start with the same level of intransparency, $\varepsilon_i = \varepsilon_j = \varepsilon$, ex ante so that $T_i > T_j$. Then*

- (i) *the stronger firm has an incentive to decrease the level of transparency;*
- (ii) *the weaker firm has no incentive to change the level of transparency.*

The stronger firm wants to discourage the weaker firm from investing by increasing uncertainty and thereby decreasing the expected profit of R&D. As the R&D investment of the firms is monotonic w.r.t. ε , the stronger firm would choose the highest possible level of intransparency. If increasing intransparency does not incur any costs, the stronger firm prefers high levels of intransparency, while the weaker firm is indifferent towards the level of intransparency.¹⁶ Hence, $\varepsilon_i^* = 1$, while $\varepsilon_j^* \in [0, 1]$ constitutes an equilibrium of

¹⁵We abstract here from the fact that changing the level of transparency might be associated with costs as, for instance, a reduced probability of obtaining a patent.

¹⁶If the ex ante level of intransparency is not the same for both firms, $T_j > T_i$ is possible. Then the weaker firm has an incentive to increase the level of intransparency, while the stronger firm is indifferent.

the game in which ε_n is chosen ex ante, if there are no bounds on the maximal level of intransparency. Note that after a discovery firms are indifferent whether to reveal their success to the competitor.

Is welfare maximal under full transparency?

Taking the benefits to consumers as well as the costs of R&D into account, the optimal level of transparency from the perspective of a welfare maximizing social planner (e.g., the patent office) can be determined. We assume that the patent office commits to a certain level of intransparency ε ex ante and that this level is the same for both firms. The patent office aims to maximize welfare, which is given by

$$\begin{aligned} \Omega = & p\psi(1 - e^{-\lambda_i T_i - \lambda_j T_j}) - c(1 - p)(T_i + T_j) - cp\varepsilon \left(\frac{1 - e^{-\lambda_j T_j}}{\lambda_j} + \frac{1 - e^{-\lambda_i T_i}}{\lambda_i} \right) \\ & - cp \left(2(1 - \varepsilon) \frac{1 - e^{-\Lambda T_j}}{\Lambda} + (\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) \frac{e^{-\lambda_i T_i} - e^{-\lambda_j T_j}}{\lambda_i} \right). \end{aligned} \quad (3.15)$$

Welfare Ω equals the sum of the firms individual expected equilibrium payoffs, where profit π is replaced by ψ , the value of a discovery for consumers. The welfare-maximizing level of transparency depends on the difference between the costs of R&D and consumer surplus, and the difference in the firms' R&D productivities. In particular, if firms are sufficiently heterogeneous in terms of their R&D productivities, positive levels of intransparency can be optimal.

Proposition 3.3. *If the patent office aims to maximize welfare by committing to the level of intransparency ε ex ante, then in a winner-takes-all competition with one patent,*

- (i) *full transparency ($\varepsilon^* = 0$) is optimal for $\partial T_i / \partial \varepsilon < 0$ and $\partial \Omega / \partial T_j > 0$;*
- (ii) *positive levels of intransparency ($\varepsilon^* > 0$) are optimal for $\partial T_i / \partial \varepsilon > 0$ while $\partial \Omega / \partial T_j < 0$, if the costs of R&D are sufficiently low.*

By analyzing the derivative of the welfare function w.r.t. ε , we can determine when full transparency maximizes welfare. This derivative has the following form

$$\frac{\partial \Omega}{\partial \varepsilon} = \frac{\partial \Omega}{\partial T_i} \frac{\partial T_i}{\partial \varepsilon} + \frac{\partial \Omega}{\partial T_j} \frac{\partial T_j}{\partial \varepsilon} + k. \quad (3.16)$$

The last term k represents the direct increase in costs that occurs if ε increases and is always negative. Moreover, welfare is increasing in the stopping time of the stronger firm

(i.e. $\partial\Omega/\partial T_i > 0$), and the stopping time of the weaker firm is decreasing in ε ($\partial T_j/\partial\varepsilon < 0$). If $\partial T_i/\partial\varepsilon < 0$ while $\partial\Omega/\partial T_j > 0$, the derivative is negative for all $\varepsilon \in [0, 1]$, welfare is monotonically decreasing in ε and $\varepsilon^* = 0$. From Proposition 3.1 we know that $\partial T_i/\partial\varepsilon < 0$ if the two firms are roughly equally strong in terms of their R&D productivities λ_i and λ_j . Now let us see under which condition welfare is increasing in the stopping time of the weaker firm.

$$\frac{\partial\Omega}{\partial T_j} > 0 \Leftrightarrow \lambda_i(\psi - \pi)e^{-\lambda_i T_i} > c(1 - \varepsilon)(e^{-\lambda_i T_j} - e^{-\lambda_i T_i}).$$

This inequality is satisfied if ψ is sufficiently large compared to c , if ε is high or if T_i is close to T_j . Hence, if firms are relatively equally strong and have similar stopping times, then welfare is maximal under full transparency. Yet, if the stronger firm is considerably more likely make a discovery and also invests more in R&D in equilibrium, $\partial T_i/\partial\varepsilon > 0$ while $\partial\Omega/\partial T_j < 0$ and the first two terms on the r.h.s. of (3.16) are positive. Hence, depending on the parameters of the model, it is possible that welfare is not necessarily maximal under full transparency, particularly if the costs of R&D are small (as this implies that k is small). If firms are sufficiently heterogeneous, then by increasing ε the R&D investment of the weaker firm can be reduced, while the R&D investment of the stronger firm increases. This increases welfare if the social planner wants to replace some of firm j 's investment by the investment of the stronger firm i . Firms that are similar in terms of their research productivities both decrease their R&D investment if ε increases and the risk of not observing a discovery of the competitor is higher.

3.5.2 Perfect R&D spillovers

Up to now we analyzed how the level of transparency affects R&D investment when firms compete for a patent in a winner-takes-all competition. However, the invention of one firm often creates positive R&D spillovers that also benefit other firms.¹⁷ In this section we analyze the impact of transparency on R&D investment in the presence of positive R&D spillovers but retain the assumption that the game ends after the first patent.

¹⁷See, e.g., Griliches (1992) or Henderson and Cockburn (1998).

Full transparency

For now suppose $\varepsilon = 0$, and $\pi_1 = \omega_1$, i.e., we have perfect positive spillovers. Let π again denote the profit from a discovery, i.e., $\pi = \pi_1\tau = \omega_1\tau$. The expected payoff of firm n is given by

$$U_n = p(\Lambda\pi - c)\frac{1 - e^{-\Lambda T-n}}{\Lambda} + p\pi e^{-\lambda_n T_n}(e^{-\lambda_{-n} T_{\min}} - e^{-\lambda_{-n} T-n}) + p(\lambda_n\pi - c)e^{-\lambda_{-n} T-n}\frac{e^{-\lambda_n T_{\min}} - e^{-\lambda_n T_n}}{\lambda_n} - (1-p)cT_n. \quad (3.17)$$

Now both firms can profit from one invention. Thus, the firm stopping earlier can still obtain a positive payoff (as an imitator) after she stopped investing, if the other firm makes a discovery. This is captured through the second term on the r.h.s. of Equation (3.17). Furthermore, while both firms invest, the instantaneous probability of a discovery is Λ as an own success is not necessary to generate profits. Taking the derivative of U_i and U_j w.r.t. T_i and T_j respectively, the first order conditions for perfect spillovers are

$$-c(1-p + pe^{-\lambda_i T_i - \lambda_j T_{\min}}) + p\pi\lambda_i e^{-\lambda_i T_i - \lambda_j T_j} = 0, \quad (3.18)$$

and

$$-c(1-p + pe^{-\lambda_i T_{\min} - \lambda_j T_j}) + p\pi\lambda_j e^{-\lambda_i T_i - \lambda_j T_j} = 0. \quad (3.19)$$

From Equations (3.18) and (3.19) the optimal stopping times in equilibrium can be derived. In this equilibrium the weaker firm experiments longer, i.e., $T_j > T_i$.

Lemma 3.3. *The optimal stopping times of firms in equilibrium for perfect R&D spillovers and full transparency ($\varepsilon = 0$) with one patent are given by*

$$T_i = -\frac{1}{\Lambda} \left[\ln\left(\frac{1-p}{p}\right) + \ln\left(\frac{\lambda_i\pi}{\lambda_j\pi - c} - 1\right) \right],$$

and

$$\begin{aligned} T_j &= -\frac{1}{\lambda_j} \left[\ln\left(\frac{1-p}{p}\right) + \ln\left(\frac{c}{\lambda_j\pi - c}\right) + \lambda_i T_i \right] \\ &= -\frac{1}{\Lambda} \left[\ln\left(\frac{1-p}{p}\right) - \ln(\lambda_j\pi - c) \right] - \frac{1}{\lambda_j} \ln(c) + \frac{\lambda_i}{\lambda_j\Lambda} \ln(\pi(\lambda_i - \lambda_j) - c). \end{aligned}$$

In this equilibrium

- (i) the weaker firm invests more, i.e., $T_j > T_i$;

(ii) the firms learn less than under cooperation, i.e., their final posterior belief in the absence of a success is more optimistic as $\lambda_i T_i + \lambda_j T_j < \lambda_i T_i^a + \lambda_j T_j^a$.

In a winner-takes-all race the stronger firm invests more, while for positive spillovers, the weaker firm invests more in R&D in a transparent patent system. The reason for this is the following. The probability of a discovery is higher for the stronger firm and so are the expected benefits of R&D. This implies that at the time of stopping the probability of paying the costs c (and thus the probability of not having observed a success so far) has to be lower for the weaker firm j , which implies that T_j has to be higher than T_i .

Also for positive spillovers, in a research alliance only the stronger firm invests in R&D. Hence, similar to the winner-takes-all competition firm j invests too much (now even more than firm i), and firm i free-rides on j 's research efforts when firms compete for patents. For positive spillovers this free-riding is so severe that firms learn less (meaning their final posterior at the time of stopping is more optimistic) than under cooperation.

Imperfect Transparency

Now suppose $\varepsilon > 0$. Then the expected payoff of firm n is

$$U_n = p\pi((1 - \varepsilon)(1 - e^{-\lambda_n T_n - \lambda_n T_n}) + \varepsilon(1 - e^{-\lambda_n T_n})) - (1 - p)cT_n - pc \left(\varepsilon \frac{1 - e^{-\lambda_n T_n}}{\lambda_n} + (1 - \varepsilon) \left(\frac{1 - e^{-\lambda T_n}}{\Lambda} + e^{-\lambda_n T_n} \frac{e^{-\lambda_n T_{\min}} - e^{-\lambda_n T_n}}{\lambda_n} \right) \right) \quad (3.20)$$

In contrast to the winner-takes-all competition, a firm can profit from her discovery even if the competitor had an unobserved success before and from a success of the competitor after stopping herself. Hence, the optimal stopping times of firms in equilibrium have to satisfy

$$-c(1 - p + pe^{-\lambda_i T_i}(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_{\min}})) + p\pi\lambda_i e^{-\lambda_i T_i}(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) = 0, \quad (F^3)$$

and

$$-c(1 - p + pe^{-\lambda_j T_j}(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_{\min}})) + p\pi\lambda_j e^{-\lambda_j T_j}(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}) = 0. \quad (F^4)$$

Equations (F^3) and (F^4) together with U_n for $n = i, j$ can be used to analyze the equilibria of the stopping game. For perfect positive R&D spillovers there does not exist a (unique)

equilibrium for all levels of transparency. To be more precise, for ε sufficiently low and sufficiently high there always exists a unique equilibrium, while for intermediate values of ε two equilibria in pure strategies or no equilibrium exist.

Proposition 3.4. *For perfect positive spillovers and transparency level $\varepsilon \in [0, 1]$,*

- (i) *the optimal stopping times of firms in equilibrium are determined by equations (F³) and (F⁴). There exist $\underline{\varepsilon}$ and $\bar{\varepsilon}$ such that for all $\varepsilon < \underline{\varepsilon}$ and for all $\varepsilon > \bar{\varepsilon}$ there exists a unique pure strategy equilibrium, where*

$$\underline{\varepsilon} = \frac{c}{\lambda_i \pi} \frac{\lambda_i \pi p - c - (\lambda_i - \lambda_j) \pi}{\lambda_j \pi p - c} \text{ and } \bar{\varepsilon} = \frac{c}{\lambda_j \pi} \frac{\lambda_j \pi p - c + (\lambda_i - \lambda_j) \pi}{\lambda_i \pi p - c}.$$

For $\varepsilon \in [\underline{\varepsilon}, \bar{\varepsilon}]$ there exist two pure strategy equilibria if and only if there exists T_n such that

$$\frac{1-p}{\varepsilon p} \left(\frac{c}{\lambda_{-n} \pi - c \varepsilon + (1-\varepsilon)e^{-\lambda_n T_n}} - \frac{c}{\lambda_n \pi e^{-\lambda_n T_n}} \right) + \frac{\lambda_n \pi - c}{\lambda_n \pi} \geq \frac{c(1-\varepsilon)}{\lambda_n \pi \varepsilon} e^{-\lambda_{-n} T_n} \quad (3.21)$$

for at least one $n = i, j$, otherwise there does not exist a pure-strategy equilibrium such that (F³) and (F⁴) are satisfied;

- (ii) *the R&D investment of firms in equilibrium is non-monotonic w.r.t. the level of intransparency ε . More precisely, there exist $\hat{\varepsilon}_n < 1$ for $n = i, j$ and $\tilde{\varepsilon} \in (0, 1)$ such that $\partial T_n(\hat{\varepsilon}_n)/\partial \varepsilon = 0$, $\lim_{\varepsilon \rightarrow \tilde{\varepsilon}} |\partial T_n(\tilde{\varepsilon})/\partial \varepsilon| = \infty$ and*

$$\frac{\partial T_n}{\partial \varepsilon} = \begin{cases} > 0 & \text{for } \varepsilon < \min\{\hat{\varepsilon}_n, \tilde{\varepsilon}\}, \\ < 0 & \text{for } \min\{\hat{\varepsilon}_n, \tilde{\varepsilon}\} < \varepsilon < \max\{\hat{\varepsilon}_n, \tilde{\varepsilon}\}, \\ > 0 & \text{for } \max\{\hat{\varepsilon}_n, \tilde{\varepsilon}\} < \varepsilon. \end{cases}$$

For perfect positive R&D spillovers there does not exist a (unique) equilibrium for all levels of transparency. For $\varepsilon < \underline{\varepsilon}$ the equilibrium is unique and the weaker firm invests more in R&D, while for $\varepsilon > \bar{\varepsilon}$ depending on the parameter values the stronger or the weaker firm invests more in the unique equilibrium. Note that $\underline{\varepsilon} = \bar{\varepsilon}$ for $\lambda_i = \lambda_j$ and the difference between the two threshold values is increasing in the difference of the firms' R&D productivities. In the interval $[\underline{\varepsilon}, \bar{\varepsilon}]$ there exist two equilibria if (3.21) is satisfied and zero otherwise. To be more precise, if (3.21) fails each firm wants to invest in R&D if the competitor does not invest and vice versa. If firm n chooses $T_n = 0$ while the competitor plays $T_{-n} = T_{-n}^m$, this pair of stopping times constitutes an equilibrium. As

firms cannot observe each others actions, they cannot coordinate on who invests and hence, there might be no equilibrium in pure strategies (because both or none of them invests). At $\varepsilon = 1$ each firm chooses T_n^m (the stopping time of a monopolist), because they cannot learn anything from the competitor. This is the highest possible level of R&D investment, as there is no combination of ε and T_{-n} such that firm n invests more than T_n^m .

The firms' R&D investment is non-monotonic in the level of transparency. Changes in the level of transparency have a direct effect on expected profits and an indirect effect through a change in the stopping time of the competitor. A higher level of intransparency reduces the probability of observing a success from the competitor and this causes firms to increase their own R&D investment. However, if the competitor increases her R&D investment, this increases the probability of observing a success from her and hence each firm has an incentive to decrease her own investment. As firms are heterogeneous in terms of their R&D productivities, the magnitude of these effects is different for each firm.

Do firms have an incentive to influence the level of transparency?

As before we are interested in whether firms have an incentive to influence the level of transparency faced by their rival. The positive spillover allows firms to profit from a discovery of the competitor. Consequently expected payoffs are increasing in the stopping time of the competitor and each firm is better off the more the competitor invests.

Proposition 3.5. *Suppose both firms start with the same level of transparency, $\varepsilon_i = \varepsilon_j = \varepsilon$, and only one patent is granted. For perfect positive spillovers each firm $n = i, j$ has an incentive to influence the level of transparency faced by the competitor to maximize the stopping time of the opponent as $\partial U_i / \partial T_j > 0$ and $\partial U_j / \partial T_i > 0$ for $n = i, j$.*

Now both firms (i.e., also firm j) have an incentive to influence the stopping time of the competitor. In contrast to the winner-takes-all competition where the stronger firm wanted to discourage the R&D investment of the weaker firm, for positive spillovers both firms want to encourage their competitor to invest. Hence, the level of transparency is chosen so as to maximize the R&D investment of the competitor. If both firms choose $\varepsilon_n^* = 1$, then in the resulting equilibrium both firms invest as much as in autarky. This is

the highest possible level of R&D investment that can be attained and none of the firms can observe a patent of the competitor. This also constitutes an equilibrium in terms of the level of intransparency as none of the firms has an incentive to change to a lower level of intransparency given the choice of the opponent. In fact, if firm n chooses $\varepsilon_n^* = 1$, firm $-n$ is indifferent towards the level of transparency faced by firm n . Hence, every combination of ε_i and ε_j in which at least one of them chooses $\varepsilon_n = 1$ is an equilibrium of the stopping game in which ε_n is chosen ex ante. As it is not clear what happens if ε is between $\underline{\varepsilon}$ and $\bar{\varepsilon}$ we assume that firms do not choose a level of transparency in this region. This is w.l.o.g. as it is always weakly optimal for each firm to choose $\varepsilon_n^* = 1$ given that the R&D investment of the competitor is maximal at $\varepsilon = 1$.

As described in Proposition 3.4, the R&D investment of the firms is non-monotonic w.r.t. the level of transparency. Hence, if firms cannot choose $\varepsilon_n = 1$, because there are bounds on the maximum level of intransparency, it is not clear whether firms would choose this upper bound, as the R&D investment of the competitor might be higher at lower levels of intransparency. The main difference to the winner-takes-all competition is thus that both firms want to encourage each other to invest and that R&D investment is non-monotonic in the level of transparency.

Is welfare maximal under full transparency?

Now we again turn to the optimal information disclosure policy of the patent office. As we assume that the value of a discovery for consumer has to be at least as high as the profit both firms can earn jointly, i.e., $\psi \geq 2\pi$, welfare for perfect spillovers is given by

$$\begin{aligned} \Omega = & p\psi(1 - e^{-\lambda_n T_n - \lambda_{-n} T_{-n}}) - c(1 - p)(T_n + T_{-n}) - cp \left(\varepsilon \frac{1 - e^{-\lambda_{-n} T_n}}{\lambda_{-n}} + \varepsilon \frac{1 - e^{-\lambda_n T_n}}{\lambda_n} \right) \\ & - cp \left(2(1 - \varepsilon) \frac{1 - e^{-\Lambda T_n}}{\Lambda} + (\varepsilon + (1 - \varepsilon)e^{-\lambda_n T_n}) \frac{e^{-\lambda_{-n} T_n} - e^{-\lambda_{-n} T_{-n}}}{\lambda_{-n}} \right), \end{aligned} \quad (3.22)$$

for $T_n < T_{-n}$. This is the same as in a winner-takes-all competition except that now it is not necessarily true that $T_i > T_j$. The welfare-maximizing level of transparency depends on the difference between the value of a discovery for consumers ψ and the costs of R&D and can be positive if the consumer surplus is sufficiently large.

Proposition 3.6. *If the patent office aims to maximize welfare by committing to the level of intransparency ε^* ex ante, for positive spillovers welfare is not necessarily maximal under full trans-*

parency. If consumer surplus is sufficiently large, i.e., for ψ sufficiently large compared to c , welfare is increasing in the R&D investment of both firms ($\partial\Omega/\partial T_j > 0$ and $\partial\Omega/\partial T_i > 0$), and high levels of intransparency ($\varepsilon^* = 1$) are optimal.

Suppose $T_j > T_i$.¹⁸ Then the derivative of the welfare function w.r.t. ε is given by

$$\begin{aligned} \frac{\partial\Omega}{\partial\varepsilon} &= \frac{\partial T_i}{\partial\varepsilon} \lambda_i e^{-\lambda_i T_i} \left(\psi e^{-\lambda_j T_j} - \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) + \frac{c(1 - \varepsilon)(e^{-\lambda_j T_i} - e^{-\lambda_j T_j})}{\lambda_j} \right) + \\ &\quad \frac{\partial T_j}{\partial\varepsilon} \lambda_j e^{-\lambda_j T_j} (\psi e^{-\lambda_i T_i} - \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i})) + k. \end{aligned}$$

As before (in the winner-takes-all competition), the last term k is negative and represents the increase in costs associated with an increase in ε . Now $\partial\Omega/\partial T_j > 0$ if

$$\psi e^{-\lambda_i T_i} > \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}),$$

which is satisfied if ψ is sufficiently large compared to π , ε is small or T_i small. Thus, it may depend on the stopping time of firm i whether welfare is increasing or decreasing in the stopping time of firm j . If it is likely that firm i does not make a discovery until reaching her stopping time, welfare is increasing in the R&D investment of firm j . If firm i is likely to make a discovery, welfare can be decreasing in T_j . Furthermore, welfare is increasing in T_i , if

$$\psi e^{-\lambda_j T_j} - \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) + \frac{c(1 - \varepsilon)(e^{-\lambda_j T_i} - e^{-\lambda_j T_j})}{\lambda_j} > 0,$$

where the last term represents the costs that are saved if firm i invests more and thereby replaces the R&D investment of firm j . For ψ sufficiently high welfare is increasing in T_i and T_j (for $T_i > T_j$ and $T_j > T_i$). Therefore, if consumer surplus is high, the patent office might prefer high levels of intransparency to encourage both firms simultaneously to invest in R&D.

The situation differs from the previous section (i.e., the winner-takes-all competition) in the following ways. In a winner-takes-all competition for high levels of ψ welfare is increasing in T_i and T_j . However, as T_j is monotonically decreasing in ε it was not clear whether the negative effect of ε on T_j and c is outweighed by an increase in T_i . For positive spillovers a higher ψ means that both firms should invest more in R&D. Furthermore,

¹⁸The corresponding expressions for $T_i > T_j$ can be found in the appendix.

	-n R&D	-n no R&D
n R&D	$\frac{\tilde{\beta}(\lambda_n\pi + \lambda_{-n}\omega - (1 - e^{-\Lambda\tau})\lambda_{-n}\alpha\tilde{\pi}) - c}{\Lambda}, \frac{\tilde{\beta}(\lambda_{-n}\pi + \lambda_n\omega) - c}{\Lambda}$	$\frac{\tilde{\beta}\lambda_n\pi - c}{\lambda_n}, \tilde{\beta}(1 - \varepsilon)\omega$
n no R&D	$\tilde{\beta}((1 - \varepsilon)\omega - (1 - e^{-\lambda_{-n}\tau})\alpha\tilde{\pi}), \frac{\tilde{\beta}\lambda_{-n}\pi - c}{\lambda_{-n}}$	0, 0

Table 3.1: Payoff matrix for R&D investment in stage 2 when firm n holds patent 1.

now welfare is not necessarily increasing in T_i (it depends j 's R&D investment and on ε). Moreover, R&D investment is non-monotonic w.r.t. ε . Hence, it is possible that there are high levels of ε for which T_i and T_j are high (both are maximal at $\varepsilon = 1$) and this also increases welfare. This is particularly likely if ψ is sufficiently large and c sufficiently small.

3.5.3 Sequential innovation

So far we analyzed the impact of transparency on strategic R&D investment for perfect spillovers as well as in a winner-takes-all competition if the game ends after the first patent is granted. Now we want to briefly discuss what changes if one innovation may lead to follow-up innovations, i.e., innovation has a sequential character and after the first patent is granted firms can invest in R&D to obtain another patent.

Stage 2: We start our analysis in the second stage and assume that firm n holds the first patent. Now it is possible that both firms are aware of the patent or that the unsuccessful firm $-n$ is unaware of n 's success. If firm $-n$ is not aware of the existing patent (meaning she is still in stage 1 or exited), she is either informed about the existing patent after an own discovery or she invests until reaching her stopping time. Suppose both firms are aware of patent 1 and decide whether to invest in a second patent. Table 1 shows the expected payoffs of firms in stage 2, where $\pi_2\tau$ is denoted by π , $\omega_2\tau$ by ω and $1 - \beta$ by $\tilde{\beta}$.

The patent holder (firm n) invests in R&D if $(1 - \beta)\lambda_n\pi_2\tau \geq c$ if firm $-n$ does not invest in R&D and if $(1 - \beta)(\lambda_n\pi_2\tau + (\varepsilon\Lambda - \lambda_n)\omega_2\tau + \alpha\tilde{\pi}_1(\Lambda(1 - e^{-\lambda_{-n}\tau}) - \lambda_{-n}(1 - e^{-\Lambda\tau}))) \geq c$ if firm $-n$ invests in R&D. For the imitator investing is optimal as long as $(1 - \beta)\lambda_{-n}\pi_2\tau \geq c$ if the inventor does not invest and if $(1 - \beta)(\lambda_{-n}\pi_2\tau + (\varepsilon\Lambda - \lambda_{-n})\omega_2\tau) \geq c$, if the inventor invests in R&D in stage 2.

If a firm invests in R&D in stage 2 the expected payoff of R&D is decreasing in β and α and does not depend on ε . The expected payoff of a firm that does not invest in stage 2 is decreasing in β and ε , and for the patent holder also in α . Thus, the expected payoff for a given strategy profile is smaller the higher the lack of clarity regarding patent scope and patent content. However, this is not true for the likelihood with that a firm invests in R&D. To be more precise, R&D investment is more likely the smaller β , the higher α , and the higher ε . The possibility that a new invention does not lead to a patent because it is seen as a duplication of an existing patent (β), reduces the probability that firms invest in R&D and the expected value of R&D. Uncertainty about the content of a patent or ex post legal disputes on the other hand increase the probability of R&D investment, but decrease the expected value of R&D. To avoid legal disputes and protect the first patent, the inventor is more likely to invest the higher α . Similarly, for positive spillovers firms are more likely to invest the higher the probability that the competitor has an unobserved patent.

Stage 1 - unclear patent scope:

First, suppose $\varepsilon = 0$, but $\alpha, \beta > 0$ and let us see how the equilibrium stopping times of the firms in stage 1 depend on these types of uncertainties. Let π_n denote $\pi_1\tau + v_n^1$ and ω_n denote $\omega_1\tau + v_n^2$, where v_n^1 denotes the expected payoff in stage 2 of the inventor and v_n^2 the expected payoff of the imitator in stage 2. For sequential innovation the stronger firm invests more in R&D (in a winner-takes-all competition and for perfect spillovers) as long as there is no uncertainty about patent content.

Lemma 3.4. *Suppose firms face a lack of transparency regarding patent scope, i.e., $\alpha, \beta > 0$, but not regarding patent content, $\varepsilon = 0$, and sequential innovations are possible. Then*

(i) *the equilibrium stopping times are given by*

$$T_i = -\frac{1}{\lambda_i} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_i \pi_i - c} \right) + \lambda_j T_j \right],$$

and

$$T_j = -\frac{1}{\Lambda} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_i \pi_i - c} \right) + \ln \left(\frac{\lambda_i \pi_i - \lambda_j \omega_j - c}{\lambda_j \pi_j - \lambda_j \omega_j - c} \right) \right];$$

(ii) *the stronger firm invests more in R&D in equilibrium, i.e., $T_i > T_j$.*

In general, it depends on the model parameters whether a welfare-maximizing social planner wants to encourage firms to invest in R&D. For now assume that consumer surplus is sufficiently high and the costs of R&D are low so that the social planner wants to maximize the probability of a discovery. The probability of a discovery (if the state of the world is good) is $1 - e^{-\lambda_i T_i - \lambda_j T_j}$ and in order to maximize this probability, $\lambda_i T_i + \lambda_j T_j$ has to be maximized. The sum $\lambda_i T_i + \lambda_j T_j$ is increasing in ω_y and π_y and also in v_n^1 and v_n^2 .

Corollary 3.1. *Suppose $\varepsilon = 0$ and a welfare-maximizing social planner aims to maximize the likelihood of an invention, $1 - e^{-\lambda_i T_i - \lambda_j T_j}$, by influencing the transparency of the patent scope, i.e., through the choice of α and β . Then*

- (i) $\alpha^* > 0$, if the inventor does not invest in stage 2 for $\alpha = 0$, but does so for $\alpha > 0$, as this increases v_n^1 and consequently $\lambda_i T_i + \lambda_j T_j$;
- (ii) $\beta^* = 0$, as v_n^1 and v_n^2 are both decreasing in β and so is $\lambda_i T_i + \lambda_j T_j$.

While the social planner prefers β to be as low as possible, meaning high levels of transparency are optimal, this is not certainly true for α , the probability of litigation. A higher risk of legal disputes can motivate the inventor to invest in follow-up innovations and thereby stimulate R&D investment.

Stage 1 - unclear patent scope and content:

Finally, let us discuss the situation where $\varepsilon > 0$ and $\alpha, \beta > 0$, i.e., there is a lack of clarity regarding patent scope and patent content. The first order conditions change to

$$e^{-\lambda_i T_i} \left(e^{-\lambda_j T_{min}} \lambda_i \pi_i + \lambda_i \omega_i (\varepsilon - e^{-\lambda_j T_{min}} + (1 - \varepsilon) e^{-\lambda_j T_j}) - c(\varepsilon + (1 - \varepsilon) e^{-\lambda_j T_{min}}) \right) = \frac{1 - p}{p} c,$$

and

$$e^{-\lambda_j T_j} \left(e^{-\lambda_i T_{min}} \lambda_j \pi_j + \lambda_j \omega_j (\varepsilon - e^{-\lambda_i T_{min}} + (1 - \varepsilon) e^{-\lambda_i T_i}) - c(\varepsilon + (1 - \varepsilon) e^{-\lambda_i T_{min}}) \right) = \frac{1 - p}{p} c,$$

These expressions can be used to analyze how the expected benefits of R&D investment vary with a lack of clarity in patent content, ε , as well as patent scope, α and β . In a winner-takes-all competition if the expected payoff of the imitator in stage 2 v_n^2 is low, (e.g., because it is likely that a new discovery will not lead to a patent, meaning β is high), the first order condition for sequential innovation is similar to a winner-takes-all competition with one patent only. If v_n^2 is high the situation is similar to the perfect positive

spillover case, meaning a winner-takes-all competition in which sequential innovation is possible is comparable to a situation with positive (but not necessarily perfect) R&D spillovers. The magnitude of v_n^2 (and particularly its difference to v_n^1) determines the level of R&D spillovers between firms. In general, $v_n^1 = v_n^2$ if firm $-n$ does not invest in R&D in stage 2 and $v_n^1 < v_n^2$ if $-n$ does invest (because of the risk of legal disputes α). However, at the same time for $\alpha > 0$, the inventor is more likely to invest. This means that it is possible that the firm invests in stage 2 as an inventor but not as an imitator and this would imply that $v_n^1 > v_n^2$. The stage 2 payoff of the imitator, v_n^2 , is decreasing in β and also in ε , yet the probability that the imitator invests is increasing in ε . Through the level of transparency the patent office can to some extent control the level of R&D spillovers. Higher levels of transparency however do not imply higher R&D spillovers. In general, it depends on the type of intransparency and also how this intransparency interacts with the spillover parameter ω_y .

In the case of sequential innovation firms also ex post have an incentive to hide an invention. By not disclosing a discovery the inventor can avoid infringement. If the inventor can not profit from a follow-up invention of the imitator (as in a winner-takes-all competition), then she will hide the first invention if possible. For positive spillovers the inventor can profit from a second discovery even if she does not achieve this discovery herself and hence, it is not clear whether she would like to hide the first patent to avoid legal disputes or disclose it to increase the probability of obtaining ω_2 .

A benevolent social planner, who wants to maximize the likelihood of an invention, prefers low levels of β . This means that a firm should be able to assess whether a new discovery will receive a patent or what steps are necessary to guarantee that it is not seen as a duplication of an existing patent. The picture is slightly more complex for the other two types of uncertainty. An inventor that has to fear ex post legal disputes is more likely to invest into improvements herself. A strong and clear patent protection (as reflected by a low α) could then cause an inventor to rest on her primary success, while a lack of clarity encourages further improvements. Nevertheless, the expected profit of R&D investment is lower if a firm expects profit reductions in the future stemming from an uncertainty of the extent of protection and this can in turn decrease the incentive to invest in R&D ex ante. Hence, the impact of α on R&D investment is not clear (and this independently

of the level of R&D spillovers). Similarly, a lack of clarity regarding patent content (ε) increases the probability that firms invest in R&D in stage 2, while the expected payoff of R&D decreases. Moreover, welfare in stage 1 is not necessarily maximal under full transparency. Particularly if consumer surplus is high and costs are low, strategic firms invest too little in R&D, because of the possibility to free-ride on the R&D investment of the competitor. In this case higher levels of intransparency can encourage R&D investment. On the other hand, if consumer surplus is low, the costs of R&D are high and the level of transparency does not alter the decision of firms whether to invest in R&D in stage 2, then welfare is increasing in the level of transparency.

3.6 Conclusion

This chapter investigated the consequences of a lack of clarity regarding the scope and the content of a patent on the incentives of firms to invest in R&D and on the optimal information disclosure policy of the patent office. Two heterogeneous firms compete for patents and strategically choose their R&D investment. In a winner-takes-all competition the R&D investment of the weaker firm is increasing in the level of transparency, while for the stronger firm this is only the case if the firms are relatively equal in terms of their R&D productivities. For positive spillovers the R&D investment of firms is non-monotonic w.r.t. to the level of transparency and maximal in a perfectly intransparent system.

R&D investment and welfare are not necessarily maximal under full transparency. If an innovation has positive R&D spillovers (either because the unsuccessful firm can easily imitate and adapt the discovery for a different market or because it may lead to follow up innovations), firms have an incentive to free-ride on each others R&D investment. This means that there is not enough duplication from the perspective of a welfare-maximizing social planner. As uncertainty about the content of a patent can lead to more duplication, it might not be optimal for the patent office to encourage full information disclosure, i.e., full transparency. Similarly, a higher risk of legal disputes does not necessarily discourage the inventor from investing. Despite the negative effect on the expected profit associated with a patent, the inventor may be more likely to invest into follow-up inventions, if she has to fear litigation.

The links between uncertainty in the patent system and strategic R&D investment are complex. While some types of uncertainty seem to have mainly a negative effect (β), others have positive as well as negative effects on R&D investment and welfare. Our theoretical findings suggest that depending on the type of uncertainty, full transparency is not necessarily optimal. Thus, increasing transparency might not only be associated with advantages, and the costs and benefits of policies to increase transparency have to be considered carefully.

During the analysis the length of a patent as well as profits π and ω were treated as exogenous and could not be influenced by the social planner. In reality policy makers can of course use these instruments, which might be controlled more easily than the transparency level itself. Ideally, the patent office will choose an optimal combination of the level of transparency, patent length and competition policies that influence ω . An interesting question for future research can thus be an analysis of the optimal combination of patent breadth, patent length and the certainty of patent protection.

Finally, we restricted attention to pure strategy equilibria, which is not entirely unproblematic as for certain parameter values there does not necessarily exist an equilibrium in pure strategies. By allowing firms to use mixed strategies or considering other strategies besides stopping strategies, we could extend our analysis and also study questions related to the timing of R&D investment or the timing of information disclosure.

Appendix A

Proof of Lemma 1.1.

Expected payoffs of the single agent are linear in effort. Thus, to calculate the optimal effort choice in $T - 1$ and T we compare expected payoffs at $U(\phi_{T-1}, \phi_T)$ for $(\phi_{T-1}, \phi_T) = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$ given that the continuation payoff $U(p_{T+1}) = 0$. That is,

$$U(0, 0) = 0,$$

$$U(1, 0) = (1 - \delta)E_{p_{T-1}} + \delta p_{T-1}E_1\pi,$$

$$U(0, 1) = \delta(1 - \delta)E_{p_{T-1}} + \delta p_{T-1}E_1\pi,$$

$$U(1, 1) = (1 - \delta)E_{p_{T-1}} + \delta E_1\pi p_{T-1}(1 + \delta - \delta\pi) + \delta(1 - \delta)(1 - p_{T-1}\pi)E_{p_T}.$$

First, $U(1, 0) \geq U(0, 0)$ for $p_{T-1} \geq p^a$. Further, $U(1, 1) \geq U(1, 0)$ for $p_T \geq p^a$ and as $p_T \geq p^a$ implies that $p_{T-1} \geq p^a$ we know that $U(1, 0) \geq U(0, 1)$. Hence, $U(1, 1) \geq U(1, 0)$ implies $U(1, 1) \geq U(0, 1)$. Finally, we can verify that $U(1, 1) > U(1, 0)$ for

$$p_{T-1} \geq \frac{(1 - \delta)|E_0|}{(1 - \delta)|E_0| + E_1(1 - \pi)(1 - \delta + \delta\pi)} = p^*,$$

where $p_T = \frac{p^*(1 - \pi)}{1 - p^*\pi} = p^a$. These arguments can be extended to show that $\phi_{T-2} = 1$ is optimal if and only if $p_{T-2} \geq p^a$ and so forth. Thus, the agent experiments if $p_t \geq p^a$ and stops otherwise. \square

Proof of Proposition 1.1.

We first derive the optimal experimentation effort in $T - 1$ and T given that there are no experiments in $T + 1$, i.e., we want to find the optimal values for (ϕ_{T-1}, ϕ_T) in a symmetric equilibrium. Linearity in the maximand implies that the solution to the maximization problem is on the boundaries of $[0, 1] \times [0, 1]$. Hence, denoting the expected payoff from action profile (ϕ_{T-1}, ϕ_T) by $U(\phi_{T-1}, \phi_T)$, we compare $U(0, 0)$, $U(1, 0)$, $U(0, 1)$ and $U(1, 1)$. Comparing expected payoffs at $\phi_{i,T-1} = 0$ and $\phi_{i,T-1} = 1$ (i.e., $U(0, 0)$ and $U(1, 0)$) we find that agent i is indifferent between experimenting and not experimenting at time $T - 1$

with no experimentation in T as long as

$$I_{T-1}^c \equiv (1 - \delta)E_{p_{T-1}} + \delta E_1 p_{T-1} \pi (1 - \phi_{j,T-1} \pi)^{n-1}$$

equals zero. If $I_{T-1}^c > 0$ the best response is to choose $\phi_{T-1}^c = 1$ and if $I_{T-1}^c < 0$ the safe arm is optimal. As players are symmetric, this is true for all agents. From I_{T-1}^c we can derive the optimal experimentation effort in a symmetric equilibrium with no experimentation in T as

$$\phi_{T-1}^c = \frac{1}{\pi} - \frac{1}{\pi} \left(\frac{(1 - \delta)|E_0|}{\delta E_1 \pi p_{T-1}} - \frac{(1 - \delta)(|E_0| + E_1)}{\delta E_1 \pi} \right)^{\frac{1}{n-1}}$$

which has $\phi_{T-1}^c = 0$ as optimal solution for beliefs p_{T-1} below p^a and $\phi_{T-1}^c = 1$ for any beliefs p_{T-1} above

$$\bar{p}^c = \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + \delta E_1 \pi (1 - \pi)^{n-1}}.$$

A similar analysis shows that $U(0, 1) > U(0, 0)$ for the same threshold beliefs p^a and \bar{p}^c for p_T . $U(1, 0)$ and $U(0, 1)$ intersect at a belief p_{T-1} which is below p^a . As $U(1, 0)$ is steeper than $U(0, 1)$ we can conclude that it is not optimal for agents to postpone experimenting. Further, $U(1, 1) \geq U(1, 0)$ for any belief

$$p_T \geq \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + E_1 \delta \pi (1 - \pi \phi_{j,T})^{n-1}},$$

which is the same belief that determines when $U(1, 0) \geq U(0, 0)$ replacing p_{T-1} by p_T and $\phi_{j,T-1}$ by $\phi_{j,T}$. Finally, $U(1, 1) \geq U(0, 0)$ for any belief p_{T-1} above

$$\frac{(1 - \delta)|E_0|(1 + \delta)}{(1 - \delta)(|E_0|(1 + \delta) + E_1) + \delta E_1 (1 - \pi \phi_{T-1}^c)^{n-1} [1 - \delta + \delta \pi (1 - \pi \phi_T^c)^{n-1}]}. \quad (\text{A.1})$$

The belief p_{T-1} at which $U(1, 1)$ intersects with $U(0, 0)$ lies above \bar{p}^c (the belief above which $U(1, 0) > U(0, 0)$) for all $\phi_{T-1}^c, \phi_T^c \in [0, 1]$. This means that if p_{T-1} is low $\phi_{i,T-1} = \phi_{i,T} = 0$ for all $i \in N$, for slightly more optimistic beliefs $\phi_{i,T-1} > 0$ while $\phi_{i,T} = 0$ and for even more optimistic belief $\phi_{i,T-1} = \phi_{i,T} = 1$. As this argument extends for any two consecutive periods in a symmetric equilibrium

$$\phi_i^c(t) = \begin{cases} 1 & \text{for } p_t \in [\bar{p}^c, 1), \\ \frac{1}{\pi} - \frac{1}{\pi} \left(\frac{(1 - \delta)|E_0|}{\delta E_1 \pi p_t} - \frac{(1 - \delta)(|E_0| + E_1)}{\delta E_1 \pi} \right)^{\frac{1}{n-1}} & \text{for } p_t \in (p^a, \bar{p}^c), \\ 0 & \text{for } p_t \in [0, p^a], \end{cases}$$

for all $i \in N$. In general, for any two consecutive periods we have either $(\phi_t^c, 0)$ or $(1, \phi_{t+1}^c)$ where $\phi_t^c, \phi_{t+1}^c \in [0, 1]$. That is, n failed experiments from \bar{p}^c yield a posterior below p^a

so that there is at most one time period in which agents do not play exclusively risky or exclusively safe. To be more precise, updating \bar{p}^c by n failed experiments yields

$$\frac{\bar{p}^c(1-\pi)^n}{1-\bar{p}^c+\bar{p}^c(1-\pi)^n} = \frac{(1-\delta)|E_0|(1-\pi)^n}{(1-\delta)|E_0|(1-\pi)^n + E_1(1-\delta+\delta\pi(1-\pi)^{n-1})}$$

which is smaller than p^a . □

Proof of Proposition 1.2.

For simplicity of exposition the main arguments are first discussed for the simpler case when $n = 4$. We proceed analogously to the proof of Proposition 1.1. Expected payoffs are linear in effort and hence to find the optimal experimentation effort (in a symmetric equilibrium) it suffices to compare the expected payoffs from experimenting with full intensity to not experimenting at all. For $n = 4$ the expected payoffs of different action profiles $U(\phi_{i,T-1}, \phi_{i,T})$ for given (symmetric) actions of the other players are

$$U(0, 0) = \delta p_{T-1} E_1 [1 - a[1 - \delta + \delta b]],$$

where $a = (1 - \phi_{j,T-2}\pi)(1 - \phi_{j,T-1}\pi)^2$ and $b = (1 - \phi_{j,T-1}\pi)(1 - \phi_{j,T}\pi)^2(1 - \delta\phi_{T,j}\pi)$

$$U(1, 0) = (1 - \delta)E_{p_{T-1}} + \delta p_{T-1} E_1 [1 - (1 - \pi)a[1 - \delta + \delta b]],$$

$$U(0, 1) = \delta(1 - \delta)E_{p_T}(1 - p_{T-1} + p_{T-1}a) + \delta p_{T-1} E_1 [1 - a[1 - \delta + \delta(1 - \pi)b]],$$

$$U(1, 1) = (1 - \delta)E_{p_{T-1}} + \delta E_1 p_{T-1} [1 - (1 - \pi)a] + \delta(1 - p_{T-1} + p_{T-1}(1 - \pi)a)[(1 - \delta)E_{p_T} + \delta E_1 [1 - (1 - \pi)b]].$$

Thus, $U(1, 0) > U(0, 0)$ for any belief p above

$$\frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + \delta\pi E_1 a [1 - \delta + \delta b]}. \quad (\text{A.2})$$

Similarly, $U(1, 0) \geq U(0, 1)$ as long as

$$p \geq \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + \delta E_1 (1 - (1 - \pi)a)}. \quad (\text{A.3})$$

It can be shown that (A.2) \geq (A.3) for all $a, b \in [0, 1]$, i.e., for all $\phi_{T-2}, \phi_{T-1}, \phi_T \in [0, 1]$. Similarly to the complete network, this implies that agents have no incentive to delay experimentation. Moreover, it can be shown that the belief at which $U(1, 1) \geq U(1, 0)$ is strictly above (A.2) and hence, there is at most one time period in which agents use interior experimentation intensities. A change in the number of players n affects expressions a and b . Independently of the number of players we have $a, b \in [0, 1]$ and hence the conclusion remains the same.

In the ring network we have to distinguish between an even and an odd number of players, as this determines how much information arrives in the last round where new information reaches agent i . As the results are similar in both cases for simplicity we will only discuss the case where n is odd here.

Let us first specify the indifference condition I_t^r for any point in time t , where up to time t all agents experimented with full intensity. The amount of information (i.e., the number of past experiments) agent i obtains in period t and subsequent periods depends on the number of agents and on how many periods already passed. At time t we have

$$I_t^r = (1 - \delta)E_{p_t} + \delta E_1 p_t \pi \left\{ (1 - \phi^r \pi)^2 (1 - \pi)^{2 \min\{t-1, d-1\}} (1 - \delta) + \delta (1 - \phi^r \pi)^4 (1 - \pi)^{4 \min\{t-1, d-1\} - 2} (1 - \delta) + \dots + \delta^{\frac{n-3}{2}} (1 - \phi^r \pi)^{n-1} (1 - \pi)^{y_{t-1}} \right\},$$

where we dropped the time index for ϕ_t^r and $y_t = (n - 3)t - 2 \sum_{x=1}^{\min\{t-1, d-1\}} (t - x)$. If $I_t^r = 0$ the agent is indifferent between experimenting and not experimenting at time t given that there are no experiments in $t + 1$. From this we can derive ϕ_t^r as well as \underline{p}_t^r and \bar{p}_t^r for any $t \geq 1$. In contrast to the complete network an explicit simple expression for ϕ_t^r cannot be derived. A proof that the root on $[0, 1]$ exists and is unique can be found below. The cutoff beliefs are increasing in t . Further, three failed experiments from \bar{p}_t^r take the posterior below the lower cutoff \underline{p}_{t+1}^r . Hence, there is at most one period in which agents use interior intensities in the symmetric equilibrium in the ring network.

Existence and uniqueness of ϕ_t^r : The expression for ϕ_1^r can be found by analyzing $I_1^r = 0$, which can be rewritten as

$$\frac{(1 - \delta)E_{p_1}}{\delta E_1 p_1 \pi} + (1 - \phi_1 \pi)^2 - [1 - (1 - \phi_1 \pi)^2] \sum_{t=1}^{\frac{n-3}{2}} \delta^t (1 - \phi_1 \pi)^{2t} = 0, \quad (\text{A.4})$$

where the expression on the l.h.s is a polynomial of order $n - 1$ in ϕ_1 . To show that the root on $[0, 1]$ is unique, it is enough to show that (A.4) is strictly monotonically decreasing for $\phi_1 \in [0, 1]$. We rewrite $I_1^r = 0$ as

$$\begin{aligned} 0 &= (1 - \phi_1 \pi)^2 - \delta (1 - \phi_1 \pi)^2 - \delta^2 (1 - \phi_1 \pi)^4 - \dots - \delta^{\frac{n-1}{2}-1} (1 - \phi_1 \pi)^{n-3} + \\ &\quad \delta (1 - \phi_1 \pi)^4 + \delta^2 (1 - \phi_1 \pi)^6 + \dots + \delta^{\frac{n-1}{2}-1} (1 - \phi_1 \pi)^{n-1} + \frac{(1 - \delta)E_{p_1}}{\delta E_1 p_1 \pi} \\ &= (1 - \delta) \left[(1 - \phi_1 \pi)^2 + \delta (1 - \phi_1 \pi)^4 + \delta^2 (1 - \phi_1 \pi)^6 + \dots + \delta^{\frac{n-1}{2}-2} (1 - \phi_1 \pi)^{n-3} \right] - \\ &\quad \delta^{\frac{n-1}{2}-1} (1 - \phi_1 \pi)^{n-1} + \frac{(1 - \delta)E_{p_1}}{\delta E_1 p_1 \pi}. \end{aligned}$$

Taking the derivative w.r.t. ϕ_1 gives

$$(1 - \delta)[-2\pi(1 - \phi_1\pi) - 4\delta\pi(1 - \phi_1\pi)^3 - \dots] - (n - 1)\delta^{\frac{n-3}{2}}\pi(1 - \phi_1\pi)^{n-2},$$

which is negative for $\phi_1 \in [0, 1]$. A similar analysis can be carried out for I_t^r for $t \geq 2$. \square

Proof of Proposition 1.3.

For simplicity we first describe the proof if agents only experiment in $t = 1$, before we extend the arguments for any $t \geq 2$. In $t = 1$ if $p_1 \geq \bar{p}^c$, all agents experiment with intensity 1 and if $p_1 \leq p^a$, no agent experiments. We want to show that for $p_1 \in (p^a, \bar{p}^c)$ equilibrium experimentation intensities are higher in the ring than in the complete network. For beliefs in (\bar{p}_1^r, \bar{p}^c) agents in the ring experiment with full intensity while players in the complete network have effort levels below 1. For beliefs in (p^a, \bar{p}_1^r) we know that in equilibrium $I_1^r = 0$ and $I_1^c = 0$. As prior beliefs are assumed to be identical it follows from $I_1^r = I_1^c$ that

$$(1 - \phi_1^c\pi)^{n-1} - (1 - \phi_1^r\pi)^2 \left(\frac{1 - \delta + \delta^{\frac{n-1}{2}}(1 - \phi_1^r\pi)^{n-3}[1 - (1 - \phi_1^r\pi)^2]}{1 - \delta(1 - \phi_1^r\pi)^2} \right) = 0. \quad (\text{A.5})$$

Equation (A.5) holds for $\phi_1^r = \phi_1^c = 0$ and in case we set $\phi_1^r = \phi_1^c = \phi_1$ the l.h.s. of (A.5) is monotonically decreasing in ϕ_1 and negative for any $\phi_1 > 0$. Consequently, for (A.5) to hold we need

$$\phi_1^r > \phi_1^c.$$

In $t \geq 2$ it has to be shown that for any prior $p_1 \geq \bar{p}^c$ (which implies $\phi_1^r = \phi_1^c = 1$), the experimentation intensity in the ring, ϕ_t^r , is at least as high as its counterpart in the complete network, ϕ_t^c . A direct comparison is not possible, as agents hold different posteriors. For ϕ_t^r and ϕ_t^c that maximize the agents' utility in the corresponding network in the interval where agents use both arms, the corresponding beliefs are given by

$$p_t^r = (1 - \delta)|E_0|/\{(1 - \delta)[|E_0| + E_1] + \delta E_1\pi[(1 - \phi^r\pi)^2(1 - \pi)^{2\min\{t-1, d-1\}}(1 - \delta) + \delta(1 - \phi^r\pi)^4(1 - \pi)^{4\min\{t-1, d-1\}-2}(1 - \delta) + \dots + \delta^{\frac{n-3}{2}}(1 - \phi^r\pi)^{n-1}(1 - \pi)^{yt-1}]\} \quad (\text{A.6})$$

for the ring and

$$p_t^c = \frac{(1 - \delta)|E_0|}{(1 - \delta)[|E_0| + E_1] + \delta E_1\pi(1 - \phi_t^c\pi)^{n-1}} \quad (\text{A.7})$$

for the complete network. Further,

$$p_t^c = \frac{p_t^r(1-\pi)^{y_{t-1}}}{p_t^r(1-\pi)^{y_{t-1}} + 1 - p_t^r}. \quad (\text{A.8})$$

By replacing p_t^r in Equation (A.8) by (A.6) and then solving for (A.8)=(A.7), it can be shown that

$$\phi_t^r > \phi_t^c.$$

□

Proof of Proposition 1.4.

Analogously to the proof of Proposition 1.1 and 1.2 it can be shown that players have no incentive to delay experimentation and that there is at most one time period in which they use interior experimentation intensities. Let us start with the central player. Comparing expected utility from experimenting with intensity ϕ_t^h to not experimenting (with no experimentation in $t+1$), the risky arm is optimal as long as $(1-\delta)E_{p_t} + \delta E_1 p_t \pi (1 - \phi_t^s \pi)^{n-1} \geq 0$. The cut-off belief above which an experimentation intensity of 1 is optimal for the hub is given by $\bar{p}^h = \bar{p}^c$, and the lower cut-off by $\underline{p}^h = p^a$. If $(1-\delta)E_{p_t} + \delta E_1 p_t \pi (1 - \phi_t^s \pi)^{n-1} = 0$, then ϕ_t^s is given by (1.8). This means that the hub is indifferent between experimenting and not experimenting on the interval $[p^a, \bar{p}^c]$ if the peripheral players choose $\phi_t^s = \phi_t^c$. If $\phi_t^s > \phi_t^c$ for a given belief then $(1-\delta)E_{p_t} + \delta E_1 p_t \pi (1 - \phi_t^s \pi)^{n-1} < 0$ and consequently the hub stops experimenting immediately. On the other hand if $\phi_t^s < \phi_t^c$, then $(1-\delta)E_{p_t} + \delta E_1 p_t \pi (1 - \phi_t^s \pi)^{n-1} > 0$ and the hub will exclusively use the risky option.

Peripheral players are symmetric and receive all their information from the hub. In $t = 1$ (with no experimentation in $t = 2$) they are indifferent between the risky and the safe arm as long as $I_1^s = 0$, where

$$I_1^s = (1-\delta)E_{p_1} + \delta E_1 p_1 \pi (1 - \phi_1^h \pi) (1 - \delta + \delta(1 - \phi_1^s \pi)^{n-2}).$$

From this we can derive ϕ_1^s and the corresponding cut-off beliefs $\underline{p}_1^s = p^a$, and

$$\bar{p}_1^s = \frac{(1-\delta)|E_0|}{(1-\delta)[|E_0| + E_1] + E_1 \delta \pi (1 - \phi_1^h \pi) (1 - \delta + \delta(1 - \pi)^{n-2})}.$$

Existence and uniqueness of ϕ_1^s can be easily verified by analyzing the expression I_1^s . The minimum of this function is at $\frac{1}{\pi}$ which implies that there is only one root on $[0, 1]$ due to the parabolic shape of the function. In $t = 1$ we have $\underline{p}^h = p^a = \underline{p}_1^s$ and $\bar{p}_1^s < \bar{p}^h = \bar{p}^c$,

where the last inequality holds for all $\phi_1^h \in [0, 1]$. Consequently, in the interval $[\bar{p}_1^s, \bar{p}^h)$, the peripheral players experiment with an effort level that violates the indifference condition of the hub ($\phi_1^s > \phi_1^c$) which implies that the central player will stop experimenting immediately for any belief below $\bar{p}^h = \bar{p}^c$. For beliefs in (p^a, \bar{p}_1^s) , if $\phi_1^h = 0$, the experimentation intensity of the peripheral players is higher than it would be in a symmetric equilibrium in the complete network. Consequently, the hub does not experiment in this region either. More precisely, comparing I_1^s and I_1^c we see that $\phi_1^s > \phi_1^c$ for all $\phi_1^h \in [0, \phi_1^c]$.

Let us turn to the problem in $t \geq 2$. The posterior beliefs of the agents are $p_{t+1}^s = \frac{p_t(1-\pi)^2}{p_t(1-\pi)^2 + 1 - p_t}$ for the peripheral players and $p_{t+1}^h = \frac{p_{t+1}^s(1-\pi)^{n-2}}{p_{t+1}^s(1-\pi)^{n-2} + 1 - p_{t+1}^s}$ for the hub. Not only do agents now hold different beliefs, also the upper and lower cut-offs for the peripheral players are different due to the first round information that will reach them. For $t \geq 2$ we have

$$I_2^s = (1 - \delta)E_{p_t^s} + \delta E_1 p_t^s \pi (1 - \pi)^{n-2} (1 - \phi_t^h \pi) (1 - \delta + \delta(1 - \phi_t^s \pi)^{n-2})$$

where, by imposing $I_2^s = 0$, we obtain

$$\bar{p}_2^s = \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + E_1 \delta \pi (1 - \pi)^{n-2} (1 - \phi_t^h \pi) (1 - \delta + \delta(1 - \pi)^{n-2})}$$

and

$$\underline{p}_2^s = \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + E_1 \delta \pi (1 - \pi)^{n-2}}.$$

We have $\bar{p}_t^h = \bar{p}^c > \underline{p}_2^s > p^a = \underline{p}^h$ and further $\bar{p}_2^s > \bar{p}^c$ for $\phi_t^h = 1$ and $t \geq 2$. Now we want to show that it is still optimal that either all agents choose effort level 1 (for high beliefs), effort level 0 (for pessimistic beliefs) or the peripheral players choose $\phi_t^s \in (0, 1)$ while the hub does not experiment. If agents in the complete network and the peripheral players have the same ϕ_t as optimal effort level, then their beliefs are less than $n - 2$ failed experiments apart from each other. This means that if the distance (measured in experiments) is $n - 2$, the belief and effort level of the peripheral players is higher than for the complete network in the interval where agents use both arms. Then it is optimal for the hub to stop experimenting completely below \bar{p}^c . As before, an optimal strategy requires either $(\phi_{i,t}^s, 0)$ or $(1, \phi_{i,t+1}^s)$, i.e., there is at most one time period in which agents use both arms simultaneously. For any belief in $(\underline{p}_2^s, \bar{p}_2^s]$, if all peripheral players choose the equilibrium experimentation effort ϕ_t^s , the posterior belief in case all experiments fail

is given by

$$\frac{(1 - \delta)|E_0|(1 - \pi\phi_t^s)(1 - \pi)^{n-2}}{(1 - \delta)|E_0|(1 - \pi\phi_t^s)(1 - \pi)^{n-2} + (1 - \delta)E_1 + E_1\delta\pi(1 - \pi)^{n-2}(1 - \delta + \delta(1 - \pi\phi_t^s)^{n-2})},$$

which is below \underline{p}_2^s . Existence and uniqueness of ϕ_t^s for $t \geq 2$ can be shown by analyzing the expression I_2^s based on the same arguments as for ϕ_1^s . \square

Proof of Proposition 1.5.

The proof consists of two parts. Part 1 is for beliefs such that in case all experiments in $t = 1$ fail, there will be no experimentation in $t = 2$. Part 2 describes the proof for beliefs where agents experiment in $t \geq 2$. First, for prior beliefs in $[0, p^a]$ and $[\bar{p}^c, 1]$ the experimentation intensity in $t = 1$ is the same in both networks. Moreover, for beliefs $\bar{p}^c - \varepsilon$ the agents in the complete network experiment almost with full intensity while the agents in the star network cannot increase their effort any more. Hence, for beliefs right below \bar{p}^c , total experimentation effort is higher in the complete network. The interesting interval is (p^a, \bar{p}_1^s) in which the hub does not experiment and in which agents both networks invest in both arms simultaneously. Therefore, $n\phi_1^c$ has to be compared to $(n - 1)\phi_1^s$. In this interval along the equilibrium path

$$p_1^c = \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + \delta E_1 \pi (1 - \phi_1^c \pi)^{n-1}},$$

for the complete network and

$$p_1^s = \frac{(1 - \delta)|E_0|}{(1 - \delta)(|E_0| + E_1) + \delta E_1 \pi (1 - \delta + \delta(1 - \phi_1^s \pi)^{n-2})},$$

for the star. This implies that for a given fixed belief the relation between ϕ_1^c and ϕ_1^s can be found through these expressions and is given by

$$\phi_1^c = \frac{1}{\pi} - \frac{1}{\pi} (1 - \delta + \delta(1 - \phi_1^s \pi)^{n-2})^{\frac{1}{n-1}}.$$

Now the difference between $\frac{n-1}{n}\phi_1^s - \phi_1^c$ can be defined as

$$\Gamma_n(\delta, \pi, p_1) := 1 - \delta - \left(1 - \frac{n-1}{n}\phi_1^s \pi\right)^{n-1} + \delta(1 - \phi_1^s \pi)^{n-2}.$$

Based on the expression for $\Gamma_n(\cdot)$ we can then define the region $S_n(p_1) \subset [0, 1]^2$ for $p_1 \in (p^a, \bar{p}^c)$ as

$$S_n(p_1) := \{\delta, \pi \in [0, 1]^2 : \Gamma_n(\delta, \pi, p_1) > 0\}.$$

That is, $S_n(p_1)$ is the set of all combinations of δ and π for which $n\phi_1^c < (n-1)\phi_1^s$. Clearly, $\Gamma_n(\delta, \pi, p_1) \rightarrow 1 - \delta > 0$ as $n \rightarrow \infty$ for all $\delta, \pi \in [0, 1]^2$ and hence $\lambda(S_n(p_1)) \rightarrow 1$ as $n \rightarrow \infty$.

If agents experiment in $t \geq 2$ a similar argument as above can be used with additionally making use of the fact that $p_t^c = \frac{p_t^s(1-\pi)^{n-2}}{p_t^s(1-\pi)^{n-2} + 1 - p_t^s}$. That is, setting

$$p_t^c = \frac{(1-\delta)|E_0|}{(1-\delta)[|E_0| + E_1] + \delta E_1 \pi (1 - \phi_t^c \pi)^{n-1}},$$

it can be solved for p_t^s , which has to be equal to

$$\frac{(1-\delta)|E_0|}{(1-\delta)[|E_0| + E_1] + \delta E_1 \pi (1 - \pi)^{n-2} (1 - \delta + \delta(1 - \phi_t^s \pi)^{n-2})}.$$

Expressing ϕ_t^c in terms of ϕ_t^s , to find out whether $\phi_t^c \leq \frac{n-1}{n}\phi_t^s$ we analyze

$$\Gamma_n(\delta, \pi, p_t) := 1 - \delta - \left(1 - \frac{n-1}{n}\phi_t^s \pi\right)^{n-1} + \delta(1 - \phi_t^s \pi)^{n-2} + \frac{(1-\delta)[1 - (1-\pi)^{n-2}]}{\delta \pi (1-\pi)^{n-2}},$$

by the same arguments as for $\Gamma_n(\delta, \pi, p_1)$. $\Gamma_n(\delta, \pi, p_t)$ is equivalent to $\Gamma_n(\delta, \pi, p_1)$ up to replacing ϕ_1^s by ϕ_t^s and adding a positive constant. $S_n(p_t) \subset [0, 1]^2$ can be defined in an analogous way as

$$S_n(p_t) = \{\delta, \pi \in [0, 1]^2 : \Gamma_n(\delta, \pi, p_t) > 0\}.$$

□

Proof of Proposition 1.6.

We start the proof by defining $F_{n,1}(\phi_1)$, which is derived by considering the difference between I_1^s and I_1^r and imposing $\phi_1^r = \phi_1^s = \phi_1$, i.e.,

$$F_{n,1}(\phi_1) := \{\delta, \pi \in [0, 1]^2 : I_1^r - I_1^s \geq 0\},$$

where $I_1^r - I_1^s$ for $\phi_1^r = \phi_1^s = \phi_1$ is given by

$$[1 - \delta(1 - \phi_1 \pi)^2][\delta - 1 - \delta(1 - \phi_1 \pi)^{n-2} + (1 - \phi_1 \pi)^2] - [1 - (1 - \phi_1 \pi)^2]\delta^{\frac{n-1}{2}}(1 - (1 - \phi_1 \pi)^{n-1}).$$

This means $F_{n,1}(\phi_1)$ represents all combinations of δ and π such that $\phi_1^s \geq \phi_1^r$. For proving part (i) of the proposition it is easy to verify that for small n (e.g., $n = 3$) the inequality is satisfied for all $\delta, \pi \in [0, 1]$. This suffices to conclude that there exists some finite $n_1 \in \mathbb{N}$ such that for all $n < n_1$ we have $F_{n,1}(\phi_1) = [0, 1]^2$. For the second part we explore the behavior of $F_{n,1}(\phi_1)$ in the limit as $n \rightarrow \infty$ and obtain

$$F_1(\phi_1) := \{\delta, \pi \in [0, 1]^2 : [1 - \delta(1 - \phi_1 \pi)^2][\delta - 1 + (1 - \phi_1 \pi)^2] \geq 0\},$$

where it can be shown that the inequality fails to hold for some values of δ and π implying that $F_1(\phi_1)$ is a strict subset of $[0, 1]^2$.

If agents experiment in $t \geq 2$ as well, we proceed in an analogous way replacing I_1^s and I_1^r with I_t^r and I_2^s and additionally make use of the fact that

$$p_t^r = \frac{p_t^s(1 - \pi)^{n-2-y_t}}{p_t^s(1 - \pi)^{n-2-y_t} + 1 - p_t^s}.$$

□

Proof of Proposition 1.7.

Part (i) is obvious as for $p_1 \in [0, p^a]$ no one experiments in any network and hence expected payoffs are zero in all networks.

To show (ii) we compare $W^c(p_1)$ with $W^s(p_1)$ making use of the fact that in the relevant interval $I_1^c = 0$ and $I_1^s = 0$ and further

$$1 - \delta + \delta(1 - \phi_1^s \pi)^{n-2} = (1 - \phi_1^c \pi)^{n-1}.$$

The result then follows from the fact that $\phi_1^s > \phi_1^c$ in equilibrium.

For (iii) we obtain the following. By comparing $W^c(p_1)$ and $W^r(p_1)$ for $p_1 \in [0, \bar{p}_1^r]$ it is straightforward to show that $c \sim r$, as in this interval $I_1^c = 0$ and $I_1^r = 0$. That $c \succ r$ for $p_1 \in (\bar{p}_1^r, 1]$ follows from discounting, i.e., the fact that $\delta < 1$.

For (iv), the condition

$$(1 - \delta)(2p_1 - 1) + \delta p_1[(1 - \pi)^{n-1}[1 + \delta(n - 1)] + (1 - \delta)(n - 1)(1 - \pi) - n(1 - \phi_1^c \pi)^n] > 0. \quad (\text{A.9})$$

is derived from $W^c(p_1) - W^s(p_1)$. That is, if (A.9) holds then $W^c(p_1) > W^s(p_1)$ for $p_1 \in (\bar{p}_1^s, \bar{p}_1^c)$.

Finally, a comparison of $W^c(p_1)$ and $W^s(p_1)$ for the case when all agents in both networks experiment with full intensity, shows that due to discounting, expected payoffs are higher in the complete network, which proves part (v). □

Proof of Proposition 1.8.

Expected payoffs in the complete network with one agent who does not experiment and no experimentation in $t \geq 2$ are given by

$$U^c = (n - 1)\delta p_1[1 - (1 - \phi^c \pi)^{n-2}] + \delta p_1[1 - (1 - \phi^c \pi)^{n-1}],$$

where we made use of the fact that for $p_1 \in [p^a, \bar{p}^c]$ we have $I_1^c = 0$. For the star network expected payoffs are

$$U^s = \delta p_1 [1 - (1 - \phi^s \pi)^{n-1}] + \delta^2 p_1 [1 - (1 - \phi^s \pi)^{n-2}] (n - 1).$$

The difference $U^c - U^s$ is given by

$$\delta p_1 \{ (1 - \phi^s \pi)^{n-1} - (1 - \phi^c \pi)^{n-1} + (n - 1) [1 - \delta + \delta (1 - \phi^s \pi)^{n-2} - (1 - \phi^c \pi)^{n-2}] \}. \quad (\text{A.10})$$

Expression (A.10) is negative, as the term in square brackets equals zero and the difference $(1 - \phi^s \pi)^{n-1} - (1 - \phi^c \pi)^{n-1}$ is negative for $\phi^s > \phi^c$.

□

Appendix B

Proof of Lemma 2.1.

At time T firm 1 compares whether investing one more time yields a higher payoff than stopping right now, i.e., when $-(1 - \delta)ck_{1,T} + \delta p_{1,T}\lambda_1 k_{1,T} \geq 0$. Linearity in $k_{1,T}$ implies that $k_{1,T} = 1$ for $p_{1,T} \geq \frac{(1-\delta)c}{\delta\lambda_1} = p_1^a$. It is straightforward to verify that $k_{1,T}^a = 1$ implies $k_{1,T-1}^a = 1$, as beliefs are decreasing and future payoffs discounted. More precisely, for any two consecutive periods we obtain

$$-k_{1,t}(1 - \delta)c + \delta p_{1,t}\lambda_1 k_{1,t} + \delta(1 - p_{1,t}\lambda_1 k_{1,t})[-k_{1,t+1}(1 - \delta)c + \delta p_{1,t+1}\lambda_1 k_{1,t+1}].$$

Now suppose $k_{1,t+1} = 1$ which implies that $p_{1,t+1} \geq p_1^a$ and $-(1 - \delta)c + \delta p_{1,t+1}\lambda_1 \geq 0$. Rearranging terms around $k_{1,t}$ to find $k_{1,t}^a$, we see that $k_{1,t}^a = 1$ if $(1 - \delta)c(\delta p_{1,t}\lambda_1 - 1) + \delta p_{1,t}\lambda_1(1 - \delta p_{1,t+1}\lambda_1) \geq 0$. This inequality is always satisfied since $p_{1,t} \geq p_{1,t+1}$. \square

Proof of Lemma 2.2.

The proof is the same as for Lemma 2.1 taking the dependency between sectors into account. At time T the firm decides whether it is better to experiment one more time compared to stopping right now, i.e., $-(1 - \delta)ck_{2,T} + \delta p_{2,T}\lambda_2 k_{2,T}[\underline{x} + p_{1,T}q_T\tilde{x}] \geq 0$. If the firm in sector 1 stopped experimenting $q_T = 0$ and we obtain the autarky cutoff $p_2^a = \frac{(1-\delta)c}{\delta\lambda_2\underline{x}}$ for sector 2. The other cutoffs can be found by distinguishing between the case when there was no breakthrough in sector 1 but firm 1 still experiments and the case when there was a breakthrough. As in Lemma 2.1, it is easy to verify that $k_{2,t+1}^a = 1$ implies $k_{2,t}^a = 1$. \square

Proof of Proposition 2.1.

For simplicity of exposition we drop the subscript indicating the sector for this proof. At time t firm $i = 1, 2$ solves the following problem:

$$u_i(p_t) = \max_{k_{i,t} \in [0,1]} \{-(1 - \delta)k_{i,t}c + \delta p_t \lambda \alpha [k_{i,t}(1 - k_{-i,t}\lambda_1(1 + \gamma)/2) + k_{-i,t}\gamma] + \delta(1 - p_t + p_t(1 - \lambda k_{i,t})(1 - \lambda k_{-i,t}))u_i(p_{t+1})\}.$$

If we substitute p_{t+1} , where

$$p_{t+1} = \frac{p_t(1 - \lambda k_{i,t})(1 - \lambda k_{-i,t})}{p_t(1 - \lambda k_{i,t})(1 - \lambda k_{-i,t}) + 1 - p_t}, \quad (\text{B.1})$$

we see that the term in brackets before $u_i(p_{t+1})$ cancels out with the denominator in (B.1) and the maximization problem is linear in $k_{i,t}$, $k_{i,t+1}$ and p_t . Similarly if we plug in for p_{t+2} we obtain linearity in $k_{i,t+2}$ and so forth. After time T agents cannot experiment anymore and $u_i(p_{T+1}) = 0$ if there was no success up to time T . Suppose $T = t + 1$ implying that $u(p_{t+2}) = 0$ and let us focus on finding the optimal values for $k_{i,t}$ and $k_{i,t+1}$. Linearity in the maximand implies that the solution to the maximization problem is on the boundaries of $[0, 1] \times [0, 1]$. Hence, denoting by $U(k_{i,t}, k_{i,t+1})$ the expected payoff of strategy profile $(k_{i,t}, k_{i,t+1})$, we need to compare the values for $U(0, 0)$, $U(1, 0)$, $U(0, 1)$, and $U(1, 1)$. These are given by

$$U(0, 0) = \delta p_t \lambda k_{-i,t} \gamma \alpha + \delta^2 p_t (1 - k_{-i,t} \lambda) k_{-i,t+1} \gamma \lambda \alpha,$$

$$U(1, 0) = -(1 - \delta)c + \delta p_t \lambda \alpha [1 - \lambda k_{-i,t} (1 + \gamma)/2 + k_{-i,t} \gamma] + \delta^2 p_t (1 - k_{-i,t} \lambda) (1 - \lambda) k_{-i,t+1} \gamma \lambda \alpha,$$

$$U(0, 1) = \delta p_t \lambda k_{-i,t} \gamma \alpha + \delta (1 - p_t k_{-i,t} \lambda) [-(1 - \delta)c + \delta p_{t+1} \lambda \alpha [1 - \lambda k_{-i,t+1} (1 + \gamma)/2 + k_{-i,t+1} \gamma]],$$

and

$$U(1, 1) = -(1 - \delta)c + \delta p_t \lambda \alpha [1 - \lambda k_{-i,t} (1 + \gamma)/2 + k_{-i,t} \gamma] - (1 - \delta) \delta c (1 - p + p(1 - \lambda)(1 - \lambda k_{i,t})) + \delta^2 p (1 - \lambda) (1 - \lambda k_{-i,t}) \lambda \alpha [1 - \lambda k_{-i,t+1} (1 + \gamma)/2 + k_{-i,t+1} \gamma].$$

Thus, $U(1, 0) \geq U(0, 0)$ if

$$p_t \geq \frac{(1 - \delta)c}{\delta \lambda \alpha [1 - \lambda k_{-i,t} (1 + \gamma)/2 - \delta \gamma (1 - k_{-i,t} \lambda) \lambda k_{-i,t+1}]} \quad (\text{B.2})$$

As players are symmetric an equivalent expression holds for player $-i$. Hence, for symmetric actions where $k_{-i,t+1} = 0$, $k_{i,t} = 0$ is a best response to $k_{-i,t} = 0$ for $p_t < \underline{p}$, both choose $k_t = 1$ for $p_t \geq \bar{p}$ and in between they choose $k_t \in (0, 1)$ so that the other firm is indifferent towards its level of R&D investment given by

$$k_t^* = \frac{-(1 - \delta)c + \delta \lambda \alpha p_t}{\delta \lambda^2 p_t \alpha (1 + \gamma)/2}.$$

A similar comparison shows that for symmetric strategies $U(0, 1) \geq U(0, 0)$, at the same threshold belief as in (B.2). Now let us compare $U(1, 0)$ and $U(0, 1)$. Both are linear and increasing in p_t and intersect with $U(0, 0)$ at the same belief for symmetric actions. As $U(1, 0)$ is steeper than $U(0, 1)$ we can conclude that $U(1, 0) \geq U(0, 1)$ for all beliefs above the intersection with $U(0, 0)$.

Further, let us see for which beliefs the action profile (1, 1) is optimal. As we know that (1, 0) dominates (0, 1) for all

$$p_t \geq \frac{(1 - \delta)c}{\delta\lambda\alpha[1 - k_t^*\lambda^{\frac{1+\gamma}{2}}]}, \quad (\text{B.3})$$

we see that $U(1, 1) - U(1, 0)$ is positive for

$$p_t \geq \frac{(1 - \delta)c}{(1 - \delta)c\lambda(2 - \lambda) + \delta\lambda(1 - \lambda)^2[1 - k_{t+1}\lambda^{\frac{1+\gamma}{2}}]}. \quad (\text{B.4})$$

If we update this to

$$p_{t+1} = \frac{p_t(1 - \lambda)^2}{p_t(1 - \lambda)^2 + 1 - p_t},$$

we obtain

$$p_{t+1} \geq \frac{(1 - \delta)c}{\delta\lambda\alpha[1 - k_{t+1}\lambda^{\frac{1+\gamma}{2}}]},$$

which is the same belief that determines when $U(1, 0) \geq U(0, 0)$ replacing k_t by k_{t+1} and p_t by p_{t+1} . Finally, we can show that $U(1, 1) \geq U(0, 0)$ for

$$p_t \geq (1 - \delta)c(1 + \delta)/\delta\lambda[[1 - \lambda k_{-i,t}(1 + \gamma)/2]\alpha + (1 - \delta)c(1 + k_{-i,t}(1 - \lambda)) + \delta\alpha(1 - k_{-i,t}\lambda)[(1 - \lambda)(1 - \lambda k_{-i,t+1}(1 + \gamma)/2) - k_{-i,t+1}\gamma\lambda]].$$

This threshold belief is above (B.3) for all $k_{-i,t}, k_{-i,t+1} \in [0, 1]$ if and only if

$$c < \frac{\delta\lambda\alpha(2 - \lambda - (1 + \gamma)/2)}{(1 - \delta)(2 - \lambda)}. \quad (\text{B.5})$$

Further, the same condition guarantees that two failed experiments from \bar{p} take the posterior belief below \underline{p} and that if both firms choose $k_{s,t}^* \in (0, 1)$ for beliefs $p_{s,t} \in (\underline{p}_s, \bar{p}_s)$, the posterior is below the lower cutoff in the absence of a success. Hence, if (B.5) is satisfied, there is at most one time period in which players use an interior action.

Note that $U(k_t, k_{t+1}, p_t)$ is linearly increasing in the belief. The threshold belief at which $U(1, 1, p_t) \geq U(1, 0, p_t)$ (given by (B.4)) is above (B.3) where $U(1, 0, p_t)$ intersects with $U(0, 0, p_t)$. Additionally, it is easy to verify that $U(1, 1)$ is steeper than $U(1, 0)$ which immediately implies that the intersection of $U(1, 1)$ with $U(0, 0)$ is above (B.3) as well. This implies that (B.5) is always satisfied. Hence, in a symmetric equilibrium for optimistic beliefs (1, k_{t+1}^*) is optimal, for low beliefs (0, 0) and for intermediate beliefs (k_t^* , 0) with $k_t^*, k_{t+1}^* \in [0, 1]$. These arguments are valid for any two consecutive periods. \square

Proof of Proposition 2.2.

A social planner maximizes aggregate expected payoffs given by

$$W(k_{i,t}, k_{-i,t}, p_t) = \max_{\{k_{i,t}, k_{-i,t}\}} \{-c(1 - \delta)(k_{i,t} + k_{-i,t}) + \delta p_{1,t} \lambda_1 (k_{i,t} + k_{-i,t} - k_{i,t} k_{-i,t} \lambda_1) W_s + \delta(1 - p_{1,t} + p_{1,t}(1 - \lambda_1 k_{i,t})(1 - \lambda_1 k_{-i,t})) W(p_{t+1})\}.$$

Setting $W(p_{t+1}) = 0$ to compare when experimenting in t is better than stopping, we obtain for symmetric actions

$$k^w = \frac{-(1 - \delta)c + \delta p_{1,t} \lambda_1 W_s}{\delta \lambda_1^2 p_{1,t} W_s},$$

which reaches its upper bound $k^w = 1$ at \bar{p}_1^w and $k^w = 0$ at \underline{p}_1^w . Similar arguments as in the proof of Proposition 2.1 show that two failed experiments from \bar{p}^w yield a posterior belief below \underline{p}^w . This implies that also for the social planner there is at most one period in which $k^w \in (0, 1)$. The second part for sector 2 can be derived in the same way. \square

Proof of Proposition 2.3.

Part (i): Firm 1 is in sector 1 while firm 2 decides whether to join or start research in (dependent) sector 2. To compare the cutoff beliefs of the strategic agent to the social planner, we separate the analysis into different intervals of p_1 .

- For $p_1 \in [\bar{p}_1, 1]$: If $p_2 < p_2^a$, sector 1 is preferred by both the social planner and the strategic firm. If $p_2 \geq p_2^a$, we have

$$u_{1,1} = -(1 - \delta)c + \delta p_1 \lambda_1 \alpha (1 + \gamma_1)(1 - \lambda_1/2) + \delta(1 - p_1 \lambda_1 (2 - \lambda_1)) u'_{1,1}, \text{ and}$$

$$u_{2,1} = -(1 - \delta)c + \delta p_2 \lambda_2 (\underline{x} + p_1 q_1 \tilde{x}) + \delta(1 - p_2 \lambda_2) u'_{2,1},$$

where $u'_{s,n}$ denotes the payoff in period $t + 1$. For the social planner we obtain

$$W_{1,1} = -2(1 - \delta)c + \delta p_1 \lambda_1 (2 - \lambda_1)(1 + \alpha) + \delta(1 - p_1 \lambda_1 (2 - \lambda_1)) W'_{1,1}, \text{ and}$$

$$W_{1,2} = -2(1 - \delta)c + \delta p_1 \lambda_1 + \delta p_2 \lambda_2 (\underline{x} + p_1 q_1 \tilde{x}) + \delta(1 - p_2 \lambda_2) u'_{2,1} + \delta(1 - p_1 \lambda_1) u'_{1,2}.$$

From these expressions we can derive the cutoff beliefs for p_2 such that for beliefs above this threshold sector 2 is preferred to sector 1. The continuation payoff in sector 2 after observing a failure, $\delta(1 - p_2 \lambda_2) u'_{2,1}$, enters both decisions in the same way.

We are only interested in the difference between the socially optimal and strategic

decision and the closed form solution of the cutoffs does not matter. Hence, we can set $u'_{2,1} = 0$ as the expression is linearly increasing in p_2 . Then we have $p_2^w > p_2^*$ if and only if

$$p_1 \lambda_1 [1 - \lambda_1 + \alpha(1 - \gamma_1)(1 - \lambda_1/2)] + (1 - p_1 \lambda_1(2 - \lambda_1))(W'_{1,1} - u'_{1,1}) - (1 - p_1 \lambda_1)u'_{1,2} > 0. \quad (\text{B.6})$$

The first term is always positive, which means that it depends on the continuation payoffs whether the inequality is satisfied. We know from Heidhues et al. (2015) that, if the optimal number of experiments for a single firm is given by K , the total number of experiments that are performed by two firms that can observe each other and interact strategically is given by $K - 1$, K or $K + 1$ (see Proposition 3 in Heidhues et al., 2015). If there are no further experiments in any sector, i.e., $u'_{1,2} = u'_{1,1} = W'_{1,1} = 0$ (which implies $K + 1$ experiments for two firms), $p_2^w > p_2^*$. What we need to show is that (B.6) is satisfied for any possible number of experiments in equilibrium. The difference between $W'_{1,1}$ and $u'_{1,1}$ is positive. Hence, it suffices verify that (B.6) holds for $W'_{1,1} = u'_{1,1} = 0$ while $u'_{1,2} > 0$. As there are at least $K - 1$ experiments if agents experiment jointly we know that for $W'_{1,1} = u'_{1,1} = 0$ in equilibrium $u'_{1,2}$ is bounded above by

$$-(1 - \delta)c + \delta p_{1,t+1} \lambda_1 + \delta(1 - p_{1,t+1} \lambda_1)[-(1 - \delta)c + \delta p_{1,t+2} \lambda_1]. \quad (\text{B.7})$$

Plugging (B.7) back in (B.6) we see that the inequality is satisfied if p_1 is smaller or equal to

$$\frac{(1 - \delta)c(1 + \delta)}{\lambda_1[(1 - \delta)c[1 + \delta(2 - \lambda_1)] - (1 - \lambda_1)[1 - \delta - \delta^2(1 - \lambda_1)] - \alpha(1 - \gamma_1)(1 - \lambda_1/2)]} := p^1.$$

Further we know that p_1 is such that the posterior belief after two failed experiments is below \underline{p}_1 , because there are no more experiments if both firms experiment one more time and fail. That is,

$$\frac{p_1(1 - \lambda_1)^2}{p_1(1 - \lambda_1)^2 + 1 - p_1} < p_1^a,$$

which implies

$$p_1 < \frac{(1 - \delta)c}{(1 - \delta)c\lambda_1(2 - \lambda_1) + \delta\lambda_1\alpha(1 - \lambda_1)^2} := p^2.$$

Now it suffices to show that $p^1 > p^2$.

- For $p_1 \in [p_1^a, \bar{p}_1)$ the firm joins sector 1 if $p_2 < p_2^a$ (the same is true for the social planner). If $p_2 \geq p_2^a$ the social planner compares the expected utility from both firms being in sector 1

$$W_{1,1} = -(1 - \delta)2k_1^*c + \delta p_1 k_1^* \lambda_1 (2 - \lambda_1 k_1^*) (1 + \alpha)$$

with expected payoffs if firms experiment in different sectors given by

$$W_{1,2} = -(1 - \delta)2c + \delta p_1 \lambda_1 \delta p_2 \lambda_2 (\underline{x} + p_1 q_1 \tilde{x}) + \delta (1 - p_2 \lambda_2) u'_{2,1} + \delta p_1 \lambda_1.$$

The strategic firm compares expected payoffs from joining sector 1 to payoffs from starting research in the dependent sector. Comparing the social planner to the strategic firm we see that $p_2^w \geq p_2^*$ if and only if

$$(1 - \delta)c(1 - k_1^*) + \delta p_1 \lambda_1 [k_1^*(2 - k_1^* \lambda_1)(1 + \alpha) - 1 - \alpha k_1^*(1 + \gamma)(1 - k_1^* \lambda_1 / 2)] \geq 0.$$

This inequality is always satisfied if k_1^* is sufficiently large, but fails to hold for small k_1^* . Thus, for p_1 close to \underline{p}_1 the social planner is less likely to let both firms experiment in sector 1.

- For $p_1 < \underline{p}_1$, firm 2 prefers sector 2 if $p_2 \geq p_2^a$, otherwise the firm is indifferent. The social planner applies the same decision rule.

The proof for the case when one firm is located in sector 2 and the other decides whether to join proceeds along the same lines. □

Proof of Proposition 2.4.

See proof of Proposition 2.3. □

Proof of Proposition 2.5.

We assume that the social planner ignores the impact of γ_s on experimentation and only cares about selecting the socially optimal sector. Expected aggregate payoffs are given by

$$W_{1,1} = -2(1 - \delta)c + \delta p_1 \lambda_1 (2 - \lambda_1) W_1 + \delta (1 - p_1 \lambda_1 (2 - \lambda_1)) W'_{1,1},$$

$$W_{1,2} = -2(1 - \delta)c + \delta p_1 \lambda_1 + \delta (1 - p_1 \lambda_1) u'_{1,2} + \delta p_2 \lambda_2 (\underline{x} + p_1 q \tilde{x}) + \delta (1 - p_2 \lambda_2) u'_{2,1}, \text{ and}$$

$$W_{2,2} = -2(1 - \delta)c + \delta p_2 \lambda_2 (2 - \lambda_2) W_2 + \delta (1 - p_2 \lambda_2 (2 - \lambda_2)) W'_{2,2}.$$

For strategic firms expected profits are given by,

$$u_{1,1} = -(1 - \delta)c + \delta p_1 \lambda_1 (1 - \lambda_1/2)(1 + \gamma_1)\alpha + \delta(1 - p_1 \lambda_1(2 - \lambda_1))u'_{1,1},$$

$$u_{2,1} = -(1 - \delta)c + \delta p_2 \lambda_2 (\underline{x} + p_1 q \tilde{x}) + \delta(1 - p_2 \lambda_2)u'_{2,1},$$

$$u_{1,2} = -(1 - \delta)c + \delta p_1 \lambda_1 + \delta(1 - p_1 \lambda_1)u'_{1,2}, \text{ and}$$

$$u_{2,2} = -(1 - \delta)c + \delta p_2 \lambda_2 (1 - \lambda_2/2)(1 + \gamma_2)\underline{x}\alpha + \delta(1 - p_2 \lambda_2(2 - \lambda_2))u'_{2,2}.$$

To find the value for γ_s at which the strategic and the socially optimal sector choice coincide, we have to find γ_1 such that $W_{1,1} - W_{1,2} = u_{1,1} - u_{2,1}$ and γ_2 such that $W_{2,2} - W_{1,2} = u_{2,2} - u_{1,2}$. By analyzing the expression for $W_{1,1} - W_{1,2} = u_{1,1} - u_{2,1}$ it is easy to see from the way \tilde{x} enters the l.h.s. and the r.h.s. that the optimal γ_1^* does not depend on this term.

From $W_{2,2} - W_{1,2} = u_{2,2} - u_{1,2}$ we obtain for $\psi_2 = (1 - \alpha\gamma_2)\underline{x}$,

$$\tilde{\gamma}_2^* = \frac{p_2 \lambda_2 [\underline{x}(1 + \alpha - \lambda_2 - \lambda_2 \alpha/2) - p_1 q_1 \tilde{x}] - (1 - p_2 \lambda_2(2 - \lambda_2))(u'_{2,2} - W'_{2,2}) - (1 - p_2 \lambda_2)u'_{2,1}}{p_2 \lambda_2 \underline{x}(1 - \lambda_2/2)\alpha}.$$

In this case \tilde{x} enters negatively (directly and indirectly through $u'_{2,1}$) and thus we have

$\tilde{\gamma}_2^* < \gamma_2^*$. If ψ_2 does not depend on γ_2 , the sign is reversed. That is, for $\psi_2 = 0$ we have

$$\tilde{\gamma}_2^* = \frac{p_2 \lambda_2 [\underline{x} + p_1 q_1 \tilde{x}] + (1 - p_2 \lambda_2(2 - \lambda_2))(u'_{2,2} - W'_{2,2}) + (1 - p_2 \lambda_2)u'_{2,1}}{p_2 \lambda_2 \underline{x}\alpha(1 - \lambda_2/2)} - 1.$$

In this case \tilde{x} enters positively (directly and indirectly through $u'_{2,1}$) and thus we have

$$\tilde{\gamma}_2^* > \gamma_2^*. \quad \square$$

Appendix C

Proof of Lemma 3.1.

The research alliance solves the following problem

$$\max_{\{T_i, T_j\}} U_i + U_j, \text{ s. t. } T_i, T_j \geq 0.$$

Using the Kuhn-Tucker approach to deal with the nonnegativity constraints, the necessary conditions for an optimum are $\frac{\partial \mathcal{L}}{\partial T_i} T_i = 0, T_i \geq 0$ while $\frac{\partial \mathcal{L}}{\partial T_i} \leq 0$ and $\frac{\partial \mathcal{L}}{\partial T_j} T_j = 0, T_j \geq 0$ while $\frac{\partial \mathcal{L}}{\partial T_j} \leq 0$. Suppose $T_i \geq T_j$. The derivatives w.r.t. T_i and T_j are

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial T_i} &= -(1-p)c + pe^{-\lambda_i T_i - \lambda_j T_j} (\lambda_i \pi - c), \text{ and} \\ \frac{\partial \mathcal{L}}{\partial T_j} &= -(1-p)c + pe^{-\lambda_i T_i - \lambda_j T_j} \frac{\lambda_i}{\lambda_j} (\lambda_i \pi - c) - pce^{-\lambda T_j} \frac{\lambda_i - \lambda_j}{\lambda_i}. \end{aligned}$$

Three possible cases have to be distinguished:

- (i) $T_i = T_j = 0$: Then $\frac{\partial \mathcal{L}}{\partial T_i} \leq 0$ and $\frac{\partial \mathcal{L}}{\partial T_j} \leq 0$, which implies

$$-c + \lambda_i \pi p \leq 0, \text{ and } -c + \lambda_j \pi p \leq 0.$$

This violates Assumption 3.1.

- (ii) $T_i, T_j > 0$: Then $\frac{\partial \mathcal{L}}{\partial T_i} = \frac{\partial \mathcal{L}}{\partial T_j} = 0$. Both derivatives cannot be zero simultaneously. If $\frac{\partial \mathcal{L}}{\partial T_i} = 0$, then $\frac{\partial \mathcal{L}}{\partial T_j} < 0$ and if $\frac{\partial \mathcal{L}}{\partial T_j} = 0$, then $\frac{\partial \mathcal{L}}{\partial T_i} > 0$. Thus, T_i and T_j cannot both be strictly positive.

- (iii) $T_i > 0, T_j = 0$: Then $\frac{\partial \mathcal{L}}{\partial T_i} = 0$ and we obtain

$$T_i = -\frac{1}{\lambda_i} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_i W^a - c} \right) \right],$$

and

$$\frac{\partial \mathcal{L}}{\partial T_j} = -c(1-p) + pe^{-\lambda_i T_i} \frac{\lambda_j}{\lambda_i} (\lambda_i \pi - c) < 0.$$

It is easy to verify that for $T_i < T_j$ there is no solution that satisfies the Kuhn-Tucker conditions. \square

Proof of Lemma 3.2.

(i) $\lambda_i > \lambda_j \implies T_i > T_j$:

Let $T_{\min} = \min\{T_i, T_j\}$, then equations (3.9) and (3.10) imply that

$$e^{-\lambda_i T_i - \lambda_j T_{\min}} (\lambda_i \pi - c) = e^{-\lambda_j T_j - \lambda_i T_{\min}} (\lambda_j \pi - c).$$

Dividing by $\lambda_j \pi - c$ and $e^{-\lambda_i T_i - \lambda_j T_{\min}}$ and taking the logarithm yields

$$\lambda_i T_i - \lambda_j T_j - (\lambda_i - \lambda_j) T_{\min} = \ln \left(\frac{\lambda_i \pi - c}{\lambda_j \pi - c} \right).$$

As the r.h.s. of this equality is positive for $\lambda_i > \lambda_j$ it follows that $\lambda_i T_i - \lambda_j T_j - (\lambda_i - \lambda_j) T_{\min}$ must be positive, which is only satisfied if $T_{\min} = T_j$.

(ii) $\lambda_i T_i + \lambda_j T_j = \lambda_i T_i^a + \lambda_j T_j^a$:

As $T_j^a = 0$, $\lambda_i T_i^a + \lambda_j T_j^a$ equals

$$-\ln \left(\frac{1-p}{p} \right) - \ln \left(\frac{c}{\lambda_i \pi - c} \right),$$

while $\lambda_i T_i + \lambda_j T_j$ is equal to

$$-\ln \left(\frac{1-p}{p} \right) - \ln \left(\frac{c}{\lambda_i \pi - c} \right) - \lambda_j T_j + \lambda_j T_j.$$

\square

Expected payoffs in a winner-takes-all competition with one patent for $\varepsilon > 0$:

The expected payoff of firm n consists of two parts: the benefit of R&D measured by the value of the patent minus the costs of R&D. The expected benefit of firm n equals

$$p\pi\lambda_n \left(\int_0^{T_{\min}} e^{-\lambda t} dt + e^{-\lambda - n T_{\min}} \int_{T_{\min}}^{T_n} e^{-\lambda n t} dt \right).$$

Firm n receives the profit π with the probability that the state of the world is good and she makes a discovery. Until T_{\min} this monopoly profit can only be obtained if none of the two firms had a discovery before. Further, if $T_n \neq T_{\min}$, firm n can still make a discovery

after firm $-n$ stopped. If $T_n = T_{\min}$ the second term in the brackets equals zero. The costs of R&D are

$$c(1-p)T_n + pc \left(\int_0^{T_{\min}} e^{-\lambda_n t} (\varepsilon + (1-\varepsilon)e^{-\lambda_{-n}t}) dt + (\varepsilon + (1-\varepsilon)e^{-\lambda_{-n}T_{\min}}) \int_{T_{\min}}^{T_n} e^{-\lambda_n t} dt \right).$$

The first term, $c(1-p)T_n$, are the costs the firm pays in case the state of the world is bad. As in this case there never will be a discovery, the firm pays until her stopping time T_n . Further, the firm pays the costs $c dt$ also in the good state of the world. The probability that n does not observe a discovery in the interval $[0, dt]$ (and hence pays the costs in the subsequent time interval) equals $e^{-\lambda_n dt} (e^{-\lambda_{-n} dt} + \varepsilon \lambda_{-n} dt)$, that is, the probability that none of the two firms has a discovery $e^{-\Lambda dt}$ or firm n did not have a discovery, while firm $-n$ had an unobserved discovery $e^{-\lambda_n dt} \varepsilon \lambda_{-n} dt$. As firm $-n$ stops to invest after a discovery, the probability of not observing a discovery from firm $-n$ in $[0, t]$ is given by $e^{-\lambda_{-n}t} + \varepsilon(1 - e^{-\lambda_{-n}t})$.

Proof of Proposition 3.1.

In part 1 we show that the stronger firm invests more in R&D for all $\varepsilon \in [0, 1]$. In part 2 it is shown that for $T_i, T_j \geq 0$ a solution to equations (F^1) and (F^2) exists and is unique. Finally, in part 3 we apply the implicit function theorem to derive $\partial T_i / \partial \varepsilon$ and $\partial T_j / \partial \varepsilon$.

Part 1: $\lambda_i > \lambda_j$ implies $T_i > T_j$ for all $\varepsilon \in [0, 1]$:

Equations (F^1) and (F^2) imply

$$c\varepsilon(e^{-\lambda_j T_j}(1 - e^{-\lambda_i T_{\min}}) - e^{-\lambda_i T_i}(1 - e^{-\lambda_j T_{\min}})) = e^{-\lambda_j T_j - \lambda_i T_{\min}}(\lambda_j \pi - c) - e^{-\lambda_i T_i - \lambda_j T_{\min}}(\lambda_i \pi - c). \quad (\text{C.1})$$

The l.h.s. of equation (C.1) is always non-negative: As $\lambda_i > \lambda_j$ it follows that $1 - e^{-\lambda_i T_{\min}} > 1 - e^{-\lambda_j T_{\min}}$. Further, rewriting (F^1) and (F^2) , yields

$$\frac{c(1-p)}{p} = e^{-\lambda_i T_i} (e^{-\lambda_j T_{\min}} (\lambda_i \pi - c(1-\varepsilon)) - c\varepsilon), \quad (\text{C.2})$$

and

$$\frac{c(1-p)}{p} = e^{-\lambda_j T_j} (e^{-\lambda_i T_{\min}} (\lambda_j \pi - c(1-\varepsilon)) - c\varepsilon). \quad (\text{C.3})$$

For $\lambda_i > \lambda_j$, $e^{-\lambda_j T_{\min}} > e^{-\lambda_i T_{\min}}$ and consequently

$$e^{-\lambda_j T_{\min}} (\lambda_i \pi - c(1-\varepsilon)) > e^{-\lambda_i T_{\min}} (\lambda_j \pi - c(1-\varepsilon)).$$

For equations (C.2) and (C.3) to be satisfied we thus need $e^{-\lambda_j T_j} > e^{-\lambda_i T_i}$, which implies $\lambda_i T_i > \lambda_j T_j$ for all $\varepsilon \in [0, 1]$. Hence, the l.h.s. of equation (C.1) is always positive and zero only if $\varepsilon = 0$. Thus, also the r.h.s. of (C.1) has to be positive, i.e.,

$$e^{-\lambda_j T_j - \lambda_i T_{\min}} (\lambda_j \pi - c) - e^{-\lambda_i T_i - \lambda_j T_{\min}} (\lambda_i \pi - c) \geq 0,$$

which can be rewritten as

$$e^{-\lambda_j T_j - \lambda_i T_{\min} + \lambda_i T_i + \lambda_j T_{\min}} \geq \frac{\lambda_i \pi - c}{\lambda_j \pi - c} > 1.$$

Taking the logarithm yields $\lambda_i T_i + \lambda_j T_{\min} - \lambda_j T_j - \lambda_i T_{\min} > 0$, which implies $T_{\min} = T_j$.

Part 2: For $T_i > T_j$ we want to show that a unique solution to the system of equations

$$-c(1 - p + pe^{-\lambda_i T_i}(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j})) + p\pi\lambda_i e^{-\lambda_i T_i - \lambda_j T_j} = 0, \quad (\text{F}^1)$$

$$-c(1 - p + pe^{-\lambda_j T_j}(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i})) + p\pi\lambda_j e^{-\lambda_j T_j - \lambda_i T_i} = 0, \quad (\text{F}^2)$$

exists. Equation (F¹) can be solved for T_i uniquely, that is,

$$T_i = -\frac{1}{\lambda_i} \left[\ln \left(\frac{1-p}{p} \right) + \ln \left(\frac{c}{\lambda_i \pi e^{-\lambda_j T_j} - c(\varepsilon + (1-\varepsilon)e^{-\lambda_j T_j})} \right) \right],$$

where $\lambda_i \pi e^{-\lambda_j T_j} - c(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) > 0$ due to (F¹). Hence it suffices to show that a solution to equation (F²) w.r.t. T_j exists and is unique. Denoting $y = e^{-\lambda_j T_j}$ equation (F²) can be rewritten as

$$\varepsilon y^{\frac{\lambda_j}{\lambda}} = -\frac{1-p}{p} + y \left(\frac{\lambda_j \pi}{c} - (1-\varepsilon) \right). \quad (\text{C.4})$$

Note that $0 \leq y \leq 1$, $\frac{\lambda_j}{\lambda} \in (0, 1)$ and we assume $\varepsilon > 0$. Thus, $f(y) = y^{\frac{\lambda_j}{\lambda}}$ is a monotonically increasing concave function going through the points (0, 0) and (1, 1). The r.h.s. of (C.4) is linear in y and has a unique intersection with $y^{\frac{\lambda_j}{\lambda}}$ on $0 \leq y \leq 1$ if and only if

$$1 \leq -\frac{1-p}{\varepsilon p} + \frac{1}{\varepsilon} \left(\frac{\lambda_j \pi}{c} - (1-\varepsilon) \right).$$

That is, at $y = 1$, the r.h.s. of (C.4) has to be greater than the l.h.s. of (C.4). This inequality is satisfied for any belief p above $\frac{c}{\lambda_j \pi}$. Hence, a solution to the system of equations exists and is unique if and only if $p \geq \frac{c}{\lambda_j \pi}$, which is always satisfied by Assumption 3.1.

Part 3: As we do not have an explicit expression for the optimal stopping times we will use the implicit function theorem (IFT) to analyze the impact of ε on the stopping times in

equilibrium. In part 2 we showed that a unique solution to (F^1) and (F^2) exists. Furthermore, both implicit functions are continuously differentiable. Let F_x^l denote the derivative of F^l w.r.t. x for $l = 1, 2$. The IFT implies that

$$\begin{bmatrix} \partial T_i / \partial \varepsilon \\ \partial T_j / \partial \varepsilon \end{bmatrix} = - \begin{bmatrix} F_{T_i}^1 & F_{T_j}^1 \\ F_{T_i}^2 & F_{T_j}^2 \end{bmatrix}^{-1} \begin{bmatrix} F_\varepsilon^1 \\ F_\varepsilon^2 \end{bmatrix}.$$

Taking the derivative of (F^1) and (F^2) w.r.t. T_i and T_j yields

$$F_{T_i}^1 = -\lambda_i p e^{-\lambda_i T_i} (e^{-\lambda_j T_j} (\lambda_i \pi - c(1 - \varepsilon)) - \varepsilon c) < 0,$$

$$F_{T_j}^1 = -\lambda_j p e^{-\lambda_i T_i - \lambda_j T_j} (\lambda_i \pi - c(1 - \varepsilon)) < 0,$$

$$F_{T_i}^2 = 0, \text{ and}$$

$$F_{T_j}^2 = -p e^{-\lambda_j T_j} (e^{-\lambda_i T_i} \Lambda (\lambda_j \pi - c(1 - \varepsilon)) - \lambda_j \varepsilon c) < 0.$$

The sign of the derivatives can be obtained by analyzing expressions (F^1) and (F^2) . From (F^1) we know that

$$p e^{-\lambda_i T_i} (e^{-\lambda_j T_j} (\lambda_i \pi - c(1 - \varepsilon)) - c\varepsilon) = (1 - p)c > 0.$$

Hence, $e^{-\lambda_j T_j} (\lambda_i \pi - c(1 - \varepsilon)) - c\varepsilon > 0$, which implies that $F_{T_i}^1 < 0$. A similar analysis can be used to verify the signs of the other derivatives. The determinant of the Jacobian is given by $F_{T_i}^1 F_{T_j}^2$, which is positive. Taking the derivative of (F^1) and (F^2) w.r.t. ε yields

$$F_\varepsilon^1 = -c p e^{-\lambda_i T_i} (1 - e^{-\lambda_j T_j}) < 0,$$

$$F_\varepsilon^2 = -c p e^{-\lambda_j T_j} (1 - e^{-\lambda_i T_i}) < 0.$$

To obtain the effect of an increase in ε on equilibrium R&D investment we need to solve

$$\begin{bmatrix} \partial T_i / \partial \varepsilon \\ \partial T_j / \partial \varepsilon \end{bmatrix} = - \begin{bmatrix} F_{T_i}^1 & F_{T_j}^1 \\ F_{T_i}^2 & F_{T_j}^2 \end{bmatrix}^{-1} \begin{bmatrix} F_\varepsilon^1 \\ F_\varepsilon^2 \end{bmatrix} = -\frac{1}{F_{T_i}^1 F_{T_j}^2} \begin{bmatrix} F_{T_j}^2 & -F_{T_j}^1 \\ 0 & F_{T_i}^1 \end{bmatrix} \begin{bmatrix} F_\varepsilon^1 \\ F_\varepsilon^2 \end{bmatrix}.$$

Thus, the derivatives can be calculated as

$$\frac{\partial T_j}{\partial \varepsilon} = -\frac{F_{T_i}^1 F_\varepsilon^2}{F_{T_i}^1 F_{T_j}^2} = -\frac{F_\varepsilon^2}{F_{T_j}^2} < 0,$$

and

$$\frac{\partial T_i}{\partial \varepsilon} = -\frac{1}{F_{T_i}^1 F_{T_j}^2} [F_{T_j}^2 F_\varepsilon^1 - F_{T_j}^1 F_\varepsilon^2],$$

where $F_{T_j}^2 F_\varepsilon^1 - F_{T_j}^1 F_\varepsilon^2$ equals

$$cp^2 e^{-\lambda_i T_i - \lambda_j T_j} \begin{bmatrix} (1 - e^{-\lambda_j T_j})(e^{-\lambda_i T_j} \Lambda(\lambda_j \pi - c + c\varepsilon) - \varepsilon c \lambda_j) \\ -\lambda_j(\lambda_i \pi - c + c\varepsilon)(1 - e^{-\lambda_i T_j})e^{-\lambda_j T_j} \end{bmatrix}.$$

Thus, $\partial T_i / \partial \varepsilon < 0$ if and only if

$$(1 - e^{-\lambda_j T_j})(e^{-\lambda_i T_j} \Lambda(\lambda_j \pi - c + c\varepsilon) - \varepsilon c \lambda_j) > \lambda_j(\lambda_i \pi - c + c\varepsilon)(1 - e^{-\lambda_i T_j})e^{-\lambda_j T_j}. \quad (\text{C.5})$$

Rewriting (C.5) by denoting λ_j by λ , λ_i by λx , where $x \geq 1$ and T_j by T , we see that $\partial T_i / \partial \varepsilon < 0$ if

$$(1 - e^{-\lambda T})(e^{-\lambda x T}(1 + x)(\lambda \pi - c + \varepsilon c) - c\varepsilon) - (\lambda x \pi - c + c\varepsilon)(1 - e^{-\lambda x T})e^{-\lambda T}$$

is positive. It is positive for $x = 1$ ($\lambda_i = \lambda_j$) and negative as x goes to ∞ . Moreover, by taking the derivative w.r.t. x we can show that the expression is decreasing in x . Hence, $\partial T_i / \partial \varepsilon > 0$ for x sufficiently large. \square

Proof of Proposition 3.2.

As U_j does not depend on T_i , firm j has no incentive to change the level of transparency. The payoff of firm i depends on the stopping time of firm j and the derivative of U_i w.r.t. T_j is

$$\frac{\partial U_i}{\partial T_j} = \frac{p\lambda_j e^{-\lambda_j T_j}}{\lambda_i} (e^{-\lambda_i T_i} - e^{-\lambda_i T_j}) [\pi \lambda_i - c(1 - \varepsilon)] < 0.$$

\square

Proof of Proposition 3.3.

The derivative of the welfare function Ω w.r.t. ε equals

$$\begin{aligned} & \frac{\partial T_i}{\partial \varepsilon} (e^{-\lambda_i T_i - \lambda_j T_j} p\psi \lambda_i - pc(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) e^{-\lambda_i T_i} - (1 - p)c) + \\ & \frac{\partial T_j}{\partial \varepsilon} \left(e^{-\lambda_i T_i - \lambda_j T_j} p\psi \lambda_j - pc \left(\varepsilon + (1 - \varepsilon) \left(e^{-\lambda_i T_j} + \frac{\lambda_j}{\lambda_i} (e^{-\lambda_i T_j} - e^{-\lambda_i T_i}) \right) \right) e^{-\lambda_j T_j} - (1 - p)c \right) \\ & + pc \left(2 \frac{1 - e^{-\lambda T_j}}{\Lambda} - \frac{1 - e^{-\lambda_j T_j}}{\lambda_j} - \frac{1 - e^{-\lambda T_j} - e^{-\lambda_i T_i} + e^{-\lambda_j T_j - \lambda_i T_i}}{\lambda_i} \right). \end{aligned}$$

Adding and subtracting $e^{-\lambda_i T_i - \lambda_j T_j} \pi p \left(\frac{\partial T_i}{\partial \varepsilon} \lambda_i + \frac{\partial T_j}{\partial \varepsilon} \lambda_j \right)$ and using equations (F^1) and (F^2) , the derivative can be rewritten as

$$\begin{aligned} \frac{\partial \Omega}{\partial \varepsilon} &= pe^{-\lambda_i T_i - \lambda_j T_j} (\psi - \pi) \left(\frac{\partial T_i}{\partial \varepsilon} \lambda_i + \frac{\partial T_j}{\partial \varepsilon} \lambda_j \right) - \frac{\partial T_j}{\partial \varepsilon} c(1 - \varepsilon)p(e^{-\lambda_i T_j} - e^{-\lambda_i T_i})e^{-\lambda_j T_j} \frac{\lambda_j}{\lambda_i} + \\ & cp \left(2 \frac{1 - e^{-\lambda T_j}}{\Lambda} - \frac{1 - e^{-\lambda_j T_j}}{\lambda_j} - \frac{1 - e^{-\lambda T_j} - e^{-\lambda_i T_i} + e^{-\lambda_j T_j - \lambda_i T_i}}{\lambda_i} \right). \end{aligned}$$

The last term represents the direct increase in costs associated with an increase in ε and is negative. Further $\partial\Omega/\partial T_i > 0$, while $\partial T_i/\partial\varepsilon$ is negative for λ_i close to λ_j . Moreover, $\partial T_j/\partial\varepsilon < 0$, while $\partial\Omega/\partial T_j$ is positive if

$$\lambda_i e^{-\lambda_i T_i} (\psi - \pi) > c(1 - \varepsilon)(e^{-\lambda_i T_j} - e^{-\lambda_i T_i}).$$

□

Proof of Lemma 3.3.

(i) $\lambda_i > \lambda_j \implies T_j > T_i$:

Let $T_{\min} = \min\{T_i, T_j\}$. Then equations (3.18) and (3.19) imply that

$$c(e^{-\lambda_i T_i - \lambda_j T_{\min}} - e^{-\lambda_j T_j - \lambda_i T_{\min}}) = \pi e^{-\lambda_i T_i - \lambda_j T_j} (\lambda_i - \lambda_j).$$

As the r.h.s. is positive for $\lambda_i > \lambda_j$ this implies that

$$-\lambda_i T_i - \lambda_j T_{\min} > -\lambda_j T_j - \lambda_i T_{\min},$$

which is only satisfied for $T_{\min} = T_i$.

(ii) $\lambda_j T_j + \lambda_i T < \lambda_i T_i^a + \lambda_j T_j^a$:

As $T_j^a = 0$, $\lambda_i T_i^a + \lambda_j T_j^a$ equals

$$-\ln\left(\frac{1-p}{p}\right) - \ln\left(\frac{c}{\lambda_i \pi^a - c}\right),$$

while $\lambda_j T_j + \lambda_i T_i$ equals

$$\lambda_i T_i - \ln\left(\frac{1-p}{p}\right) - \ln\left(\frac{c}{\lambda_j \pi - c}\right) - \lambda_i T_i,$$

which implies $T_j \lambda_j + \lambda_i T_i < \lambda_i T_i^a + \lambda_j T_j^a$ as

$$\ln(\lambda_j \pi - c) < \ln(\lambda_i \pi^a - c),$$

where $\pi^a = \pi_1 \tau + \omega_1 \tau = 2\pi > \pi$.

□

Proof of Proposition 3.4.

The proof is similar to the proof of Proposition 3.1. In part 1 conditions for existence and uniqueness are derived. In part 2 we analyze the impact of ε on the optimal stopping times.

Part 1: Existence and uniqueness of a solution to the system of equations

$$-c(1-p+pe^{-\lambda_i T_i}(\varepsilon+(1-\varepsilon)e^{-\lambda_j T_{\min}}))+p\pi\lambda_i e^{-\lambda_i T_i}(\varepsilon+(1-\varepsilon)e^{-\lambda_j T_j})=0, \quad (F^3)$$

$$-c(1-p+pe^{-\lambda_j T_j}(\varepsilon+(1-\varepsilon)e^{-\lambda_i T_{\min}}))+p\pi\lambda_j e^{-\lambda_j T_j}(\varepsilon+(1-\varepsilon)e^{-\lambda_i T_i})=0. \quad (F^4)$$

First, note that $e^{-\lambda_j T_{\min}} = \max\{e^{-\lambda_j T_i}, e^{-\lambda_j T_j}\}$ and $e^{-\lambda_i T_{\min}} = \max\{e^{-\lambda_i T_i}, e^{-\lambda_i T_j}\}$. Thus, we can rewrite (F³) and (F⁴) as

$$-c(1-p)+\varepsilon p x(\lambda_i \pi-c)+(1-\varepsilon) p x(\lambda_i \pi y-c \max\{y, x^{\mu_i}\})=0,$$

$$-c(1-p)+\varepsilon p y(\lambda_j \pi-c)+(1-\varepsilon) p y(\lambda_j \pi x-c \max\{x, y^{\mu_j}\})=0,$$

where $x = e^{-\lambda_i T_i}$, $y = e^{-\lambda_j T_j}$, $\mu_i = \frac{\lambda_j}{\lambda_i}$, $\mu_j = \frac{\lambda_i}{\lambda_j}$ and $x, y \in [0, 1]$. Now two possible cases have to be distinguished:

(i) $y > x^{\mu_i}$ which implies $e^{-\lambda_j T_j} > e^{-\lambda_j T_i}$ and hence $T_i > T_j$: Then we can solve (F³) for x given by

$$x = \frac{c(1-p)}{p(\lambda_i \pi - c)} \frac{1}{\varepsilon + (1-\varepsilon)y}, \quad (C.6)$$

which is monotonically decreasing in y for $y \in [0, 1]$. Solving (F⁴) for x yields

$$x = \frac{c(1-p)}{(1-\varepsilon)p\lambda_j\pi} \frac{1}{y} - \frac{\varepsilon(\lambda_j\pi - c)}{(1-\varepsilon)\pi\lambda_j} + \frac{c}{\pi\lambda_j} y^{\mu_j}, \quad (C.7)$$

which is convex and has one global minimum (possibly on $[0, 1]$) for $y \in [0, 1]$. The system of equations (C.6) and (C.7) has at most two solutions. At $y = 0$, (C.6) is smaller than (C.7). A necessary and sufficient condition for existence and uniqueness is that (C.6) > (C.7) at $y = 1$, that is,

$$\frac{c(1-p)}{p(\lambda_i\pi - c)} > \frac{c(1-p)}{(1-\varepsilon)p\lambda_j\pi} - \frac{\varepsilon(\lambda_j\pi - c)}{(1-\varepsilon)\pi\lambda_j} + \frac{c}{\pi\lambda_j}.$$

This is equivalent to

$$\varepsilon > \frac{c}{\lambda_j\pi} \frac{\pi(\lambda_i - \lambda_j) + \lambda_j p\pi - c}{\lambda_i p\pi - c} = \bar{\varepsilon}.$$

A necessary and sufficient condition for the system of equations (C.6) and (C.7) to have two solutions is that there exists T_i such that

$$\frac{1-p}{p} \left(\frac{c}{\lambda_j \pi - c \varepsilon + (1-\varepsilon)e^{-\lambda_i T_i}} - \frac{c}{\lambda_i \pi (1-\varepsilon)e^{-\lambda_i T_i}} \right) + \frac{\varepsilon}{1-\varepsilon} \left(1 - \frac{c}{\lambda_i \pi} \right) - \frac{c}{\lambda_i \pi} e^{-\lambda_j T_i} \geq 0.$$

(ii) $y < x^{\mu_i}$, which implies $T_j > T_i$. By similar arguments as for (i) a necessary and sufficient condition for existence and uniqueness is

$$\varepsilon < \frac{c}{\lambda_i \pi} \frac{\lambda_i p \pi - c - \pi(\lambda_i - \lambda_j)}{\lambda_j p \pi - c} = \underline{\varepsilon}.$$

Similarly as for $T_i > T_j$ a necessary and sufficient condition for the system of equations (C.6) and (C.7) to have two solutions is that there exists T_j such that

$$\frac{1-p}{p} \left(\frac{c}{\lambda_i \pi - c \varepsilon + (1-\varepsilon)e^{-\lambda_j T_j}} - \frac{c}{\lambda_j \pi (1-\varepsilon)e^{-\lambda_j T_j}} \right) + \frac{\varepsilon}{1-\varepsilon} \left(1 - \frac{c}{\lambda_j \pi} \right) - \frac{c}{\lambda_j \pi} e^{-\lambda_i T_j} \geq 0.$$

Part 2: We again use the implicit function theorem to derive $\partial T_i / \partial \varepsilon$ and $\partial T_j / \partial \varepsilon$. In part 1 we derived conditions for existence and uniqueness. Further, both equations are continuously differentiable. The derivatives of (F³) and (F⁴) for $T_j > T_i$ w.r.t. T_i and T_j are

$$\frac{\partial F^3}{\partial T_i} = -\lambda_i^2 p \pi e^{-\lambda_i T_i} (\varepsilon + (1-\varepsilon)e^{-\lambda_j T_j}) + c p e^{-\lambda_i T_i} (\lambda_i \varepsilon + (1-\varepsilon)\Lambda e^{-\lambda_j T_i}),$$

$$\frac{\partial F^3}{\partial T_j} = -p \lambda_j \lambda_i (1-\varepsilon) \pi e^{-\lambda_i T_i - \lambda_j T_j} < 0,$$

$$\frac{\partial F^4}{\partial T_i} = -p \lambda_i (1-\varepsilon) e^{-\lambda_i T_i - \lambda_j T_j} (\lambda_j \pi - c) < 0, \text{ and}$$

$$\frac{\partial F^4}{\partial T_j} = -p \lambda_j e^{-\lambda_j T_j} (\lambda_j \pi - c) (\varepsilon + (1-\varepsilon)e^{-\lambda_i T_i}) < 0.$$

The determinant of the Jacobian $|J|$ is given by $F_{T_i}^3 F_{T_j}^4 - F_{T_j}^3 F_{T_i}^4$, and equals

$$p^2 \lambda_j (\lambda_j \pi - c) e^{-\lambda_i T_i - \lambda_j T_j} \left(\begin{aligned} &\varepsilon^2 ((\lambda_i^2 \pi - \lambda_i c + \Lambda c e^{-\lambda_j T_i})(1 - e^{-\lambda_i T_i}) - \lambda_i^2 \pi e^{-\lambda_j T_j}) - \Lambda c e^{-\Lambda T_i} \\ &+ \varepsilon (e^{-\lambda_i T_i} \lambda_i (\lambda_i \pi - c) + \lambda_i^2 \pi e^{-\lambda_j T_j} - (1 - e^{-\lambda_i T_i}) \Lambda c e^{-\lambda_j T_i}) \end{aligned} \right).$$

Further, the derivatives w.r.t. ε are

$$F_\varepsilon^3 = p e^{-\lambda_i T_i} (\lambda_i \pi (1 - e^{-\lambda_j T_j}) - c (1 - e^{-\lambda_j T_i})) > 0, \text{ and}$$

$$F_\varepsilon^4 = p e^{-\lambda_j T_j} (\pi \lambda_j - c) (1 - e^{-\lambda_i T_i}) > 0.$$

Consequently,

$$\frac{\partial T_j}{\partial \varepsilon} = -\frac{1}{F_{T_i}^3 F_{T_j}^4 - F_{T_j}^3 F_{T_i}^4} (-F_{T_i}^4 F_\varepsilon^3 + F_{T_i}^3 F_\varepsilon^4),$$

and

$$\frac{\partial T_i}{\partial \varepsilon} = -\frac{1}{F_{T_i}^3 F_{T_j}^4 - F_{T_j}^3 F_{T_i}^4} (F_{T_j}^4 F_\varepsilon^3 - F_{T_j}^3 F_\varepsilon^4),$$

where $-F_{T_i}^4 F_\varepsilon^3 + F_{T_i}^3 F_\varepsilon^4$ equals

$$p^2 e^{-\lambda_j T_j - \lambda_i T_i} (\lambda_j \pi - c) \begin{pmatrix} c(\lambda_i \varepsilon - \lambda_i e^{-\lambda_i T_i} + (1 - \varepsilon)e^{-\lambda_j T_i} (\Lambda - \lambda_j e^{-\lambda_i T_i})) + \\ \lambda_i^2 \pi ((1 - \varepsilon)(e^{-\lambda_i T_i} - e^{-\lambda_j T_j}) - \varepsilon(1 - e^{-\lambda_i T_i})) \end{pmatrix}$$

and $F_{T_j}^4 F_\varepsilon^3 - F_{T_j}^3 F_\varepsilon^4$ is given by

$$p^2 e^{-\lambda_j T_j - \lambda_i T_i} \lambda_j (\lambda_j \pi - c) \begin{pmatrix} c(1 - e^{-\lambda_j T_i})(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}) + \\ \lambda_i \pi (e^{-\lambda_j T_j} - \varepsilon - e^{-\lambda_i T_i}(1 - \varepsilon)) \end{pmatrix}.$$

Thus,

$$\frac{\partial T_i}{\partial \varepsilon} = -\frac{c(1 - e^{-\lambda_j T_i})(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}) + \lambda_i \pi (e^{-\lambda_j T_j} - \varepsilon - e^{-\lambda_i T_i}(1 - \varepsilon))}{\lambda_i^2 \pi \varepsilon (1 - \varepsilon) e^{-\lambda_j T_j} + (\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i})(\lambda_i \varepsilon (\lambda_i \pi - c) - \Lambda c(1 - \varepsilon) e^{-\lambda_j T_i})},$$

$$\frac{\partial T_j}{\partial \varepsilon} = -\frac{1}{\lambda_j} \frac{\lambda_i (\varepsilon - e^{-\lambda_i T_i})(c - \lambda_i \pi) + (1 - \varepsilon)(e^{-\lambda_j T_i} \Lambda c - \lambda_i^2 \pi e^{-\lambda_j T_j} - c \lambda_j e^{-\Lambda T_i})}{\lambda_j \lambda_i^2 \pi \varepsilon (1 - \varepsilon) e^{-\lambda_j T_j} + (\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i})(\lambda_i \varepsilon (\lambda_i \pi - c) - \Lambda c(1 - \varepsilon) e^{-\lambda_j T_i})}.$$

To analyze these expressions we start by considering the denominator, a quadratic function in ε of the shape $b_0 + b_1 \varepsilon + b_2 \varepsilon^2$, where $b_0 = -c e^{-\Lambda T_i} \Lambda < 0$, $b_1 = \lambda_i^2 \pi (e^{-\lambda_i T_i} + e^{-\lambda_j T_j}) - c(\Lambda e^{-\lambda_j T_i} (1 - 2e^{-\lambda_i T_i}) + e^{-\lambda_i T_i} \lambda_i)$ and $b_2 = \lambda_i^2 \pi (1 - e^{-\lambda_i T_i} - e^{-\lambda_j T_j}) - c((1 - e^{-\lambda_i T_i}) \lambda_i - \Lambda e^{-\lambda_j T_i} (1 - e^{-\lambda_i T_i}))$. At $\varepsilon = 1$, $b_0 + b_1 + b_2 = \lambda_i (\lambda_i \pi - c) > 0$. This implies that on $\varepsilon \in [0, 1]$ there is exactly one root. Let $\tilde{\varepsilon}$ denote this root. The function $1/(b_0 + b_1 \varepsilon + b_2 \varepsilon^2)$ is negative for $\varepsilon < \tilde{\varepsilon}$ and goes to $-\infty$ as $\varepsilon \rightarrow \tilde{\varepsilon}^-$, it is positive for $\varepsilon > \tilde{\varepsilon}$ and goes to ∞ as $\varepsilon \rightarrow \tilde{\varepsilon}^+$. Now consider the nominator of $\partial T_i / \partial \varepsilon$, which is linear in ε and of the shape $a_0 + a_1 \varepsilon$, where $a_0 = \lambda_i \pi (e^{-\lambda_j T_j} - e^{-\lambda_i T_i}) + c e^{-\lambda_i T_i} (1 - e^{-\lambda_j T_i})$ and $a_1 = (1 - e^{-\lambda_i T_i})(c(1 - e^{-\lambda_j T_i}) - \lambda_i \pi) < 0$. $\partial T_i / \partial \varepsilon = 0$ at $\varepsilon = \hat{\varepsilon}_i = -a_0 / a_1 < 1$. Multiplying $(a_0 + a_1 \varepsilon)$ by $1/(b_0 + b_1 \varepsilon + b_2 \varepsilon^2)$ gives us the sign of the derivative. A similar analysis can be used to determine the sign of $\partial T_j / \partial \varepsilon$, as well as the signs of the derivatives for $T_j < T_i$. \square

Proof of Proposition 3.5.

For $T_j > T_i$, the expected payoffs are given by

$$U_j = -(1 - p)cT_j + p(\lambda_j \pi - c) \left(\varepsilon_i \frac{1 - e^{-\lambda_j T_i}}{\lambda_j} + (1 - \varepsilon_i) \frac{1 - e^{-\Lambda T_i}}{\Lambda} \right) + p(\lambda_j \pi - c)(\varepsilon_i + (1 - \varepsilon_i)e^{-\lambda_i T_i}) \frac{e^{-\lambda_j T_i} - e^{-\lambda_j T_j}}{\lambda_j}$$

and

$$U_i = -(1 - p)cT_i + p(\lambda_i \pi - c) \left(\varepsilon_j \frac{1 - e^{-\lambda_i T_i}}{\lambda_i} + (1 - \varepsilon_j) \frac{1 - e^{-\Lambda T_i}}{\Lambda} \right) + p\pi e^{-\lambda_i T_i} (e^{-\lambda_j T_i} - e^{-\lambda_j T_j}).$$

Hence,

$$\frac{\partial U_i}{\partial T_j} = p\pi\lambda_j(1 - \varepsilon_j)e^{-\lambda_i T_i - \lambda_j T_j} > 0,$$

and

$$\frac{\partial U_j}{\partial T_i} = p(1 - \varepsilon_i)\frac{\lambda_i}{\lambda_j}e^{-\lambda_i T_i}((\lambda_j\pi - c)e^{-\lambda_j T_j} + ce^{-\lambda_j T_i}) > 0.$$

Similarly, $\partial U_i/\partial T_j > 0$ and $\partial U_j/\partial T_i > 0$ for $T_i > T_j$. □

Proof of Proposition 3.6.

Taking the derivative of Ω in (3.22) w.r.t. ε yields (for $T_i < T_j$)

$$\begin{aligned} & \frac{\partial T_j}{\partial \varepsilon} (e^{-\lambda_i T_i - \lambda_j T_j} p\psi\lambda_j - pc(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}) e^{-\lambda_j T_j} - (1 - p)c) + \\ & \frac{\partial T_i}{\partial \varepsilon} \left(e^{-\lambda_i T_i - \lambda_j T_j} p\psi\lambda_j - pc \left(\varepsilon + (1 - \varepsilon)(e^{-\lambda_j T_i} - \frac{\lambda_i}{\lambda_j}(e^{-\lambda_j T_i} - e^{-\lambda_j T_j})) \right) e^{-\lambda_i T_i} - (1 - p)c \right) \\ & + pc \left(2\frac{1 - e^{-\Lambda T_i}}{\Lambda} - \frac{1 - e^{-\lambda_i T_i}}{\lambda_i} - \frac{1 - e^{-\Lambda T_i} - e^{-\lambda_j T_j} + e^{-\lambda_j T_j - \lambda_i T_i}}{\lambda_j} \right). \end{aligned}$$

Adding and subtracting $p\pi\lambda_i e^{-\lambda_i T_i}(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j})\frac{\partial T_i}{\partial \varepsilon}$ and $p\pi\lambda_j e^{-\lambda_j T_j}(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i})\frac{\partial T_j}{\partial \varepsilon}$, the derivative can be rewritten as

$$\begin{aligned} \frac{\partial \Omega}{\partial \varepsilon} &= \frac{\partial T_j}{\partial \varepsilon} \lambda_j e^{-\lambda_j T_j} (\psi e^{-\lambda_i T_i} - \pi\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}) + \\ & \frac{\partial T_i}{\partial \varepsilon} \lambda_i e^{-\lambda_i T_i} \left(\psi e^{-\lambda_j T_j} - \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) + \frac{c(1 - \varepsilon)(e^{-\lambda_j T_i} - e^{-\lambda_j T_j})}{\lambda_j} \right) + \\ & c \left(2\frac{1 - e^{-\Lambda T_i}}{\Lambda} - \frac{1 - e^{-\lambda_i T_i}}{\lambda_i} - \frac{1 - e^{-\Lambda T_i} - e^{-\lambda_j T_j} + e^{-\lambda_j T_j - \lambda_i T_i}}{\lambda_j} \right). \end{aligned}$$

The last term is negative and represents the increase in costs associated with an increase in ε . Now $\partial\Omega/\partial T_j > 0$ if

$$\psi e^{-\lambda_i T_i} > \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}).$$

Furthermore, welfare is increasing in T_i , if

$$\psi e^{-\lambda_j T_j} - \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j}) + \frac{c(1 - \varepsilon)(e^{-\lambda_j T_i} - e^{-\lambda_j T_j})}{\lambda_j} > 0.$$

For $T_i > T_j$, the derivative changes to

$$\begin{aligned} \frac{\partial \Omega}{\partial \varepsilon} &= \frac{\partial T_j}{\partial \varepsilon} \lambda_j e^{-\lambda_j T_j} \left(\psi e^{-\lambda_i T_i} - \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_i T_i}) - \frac{c(1 - \varepsilon)(e^{-\lambda_i T_j} - e^{-\lambda_i T_i})}{\lambda_i} \right) + \\ & \frac{\partial T_i}{\partial \varepsilon} \lambda_i e^{-\lambda_i T_i} (\psi e^{-\lambda_j T_j} - \pi(\varepsilon + (1 - \varepsilon)e^{-\lambda_j T_j})) + \\ & c \left(2\frac{1 - e^{-\Lambda T_i}}{\Lambda} - \frac{1 - e^{-\lambda_i T_i}}{\lambda_i} - \frac{1 - e^{-\Lambda T_i} - e^{-\lambda_j T_j} + e^{-\lambda_j T_j - \lambda_i T_i}}{\lambda_j} \right). \end{aligned}$$

□

Proof of Lemma 3.4.

The expected costs of R&D do not change compared to the previous sections, while the expected benefit for firm n is given by

$$\begin{aligned} & (\pi_n \lambda_n + \omega_n \lambda_{-n}) \left(\frac{1 - e^{-\Lambda T_{\min}}}{\Lambda} \right) + e^{-\lambda_{-n} T_{-n}} \frac{e^{-\lambda_n T_{\min}} - e^{-\lambda_n T_n}}{\lambda_n} \pi_n \lambda_n \\ & + \omega_n e^{-\lambda_n T_{\min}} (e^{-\lambda_{-n} T_{\min}} - e^{-\lambda_{-n} T_{-n}}), \end{aligned}$$

where $\pi_n = \pi_1 \tau + v_n^1$ and $\omega_n = \omega_1 \tau + v_n^2$. Hence the first order conditions are given by

$$-(1-p)c + p\omega_i \lambda_i e^{-\lambda_j T_j - \lambda_i T_i} + p(\lambda_i(\pi_i - \omega_i) - c)e^{-\lambda_i T_i - \lambda_j T_{\min}} = 0, \quad (\text{C.8})$$

and

$$-(1-p)c + p\omega_j \lambda_j e^{-\lambda_i T_i - \lambda_j T_j} + p(\lambda_j(\pi_j - \omega_j) - c)e^{-\lambda_j T_j - \lambda_i T_{\min}} = 0. \quad (\text{C.9})$$

This implies

$$\begin{aligned} e^{-\lambda_i T_i - \lambda_j T_j} (\lambda_i \omega_1 + \lambda_i v_i^2 - \lambda_j \omega_1 - \lambda_j v_j^2) &= e^{-\lambda_j T_j - \lambda_i T_{\min}} (\lambda_j(\pi_j - \omega_j) - c) - \\ & e^{-\lambda_i T_i - \lambda_j T_{\min}} (\lambda_i(\pi_i - \omega_i) - c). \end{aligned}$$

For $\lambda_i > \lambda_j$ the l.h.s. is always positive no matter whether both firms invest, only the inventor invests, only the imitator invests or none of the firm invests. Thus,

$$e^{-\lambda_j T_j - \lambda_i T_{\min}} (\lambda_j(\pi_j - \omega_j) - c) > e^{-\lambda_i T_i - \lambda_j T_{\min}} (\lambda_i(\pi_i - \omega_i) - c),$$

which implies that

$$-\lambda_j T_j - \lambda_i T_{\min} + \lambda_i T_i + \lambda_j T_{\min} > \ln \left(\frac{\lambda_i(\pi_i - \omega_i) - c}{\lambda_j(\pi_j - \omega_j) - c} \right) \quad (\text{C.10})$$

The r.h.s. of (C.10) is always greater or equal to 1, which implies $T_{\min} = T_j$. □

Proof of Corollary 3.1.

From Lemma 3.2 and 3.4 we know that in a winner-takes-all competition $\lambda_i T_i + \lambda_j T_j$ equals

$$-\ln \left(\frac{1-p}{p} \right) - \ln \left(\frac{c}{\lambda_i \pi_i - c} \right),$$

which is increasing in π_i . For perfect positive spillovers $\lambda_i T_i + \lambda_j T_j$ equals

$$-\ln \left(\frac{1-p}{p} \right) - \ln \left(\frac{c}{\lambda_j \pi_j - c} \right),$$

which is as well increasing in π_j . □

Bibliography

Acemoglu, D., Bimpikis, K., Ozdaglar, A. (2011), Experimentation, patents, and innovation, *American Economic Journal: Microeconomics* 3: 37-77.

Acemoglu, D., Bimpikis, K., Ozdaglar, A. (2014), Dynamics of information exchange in endogenous social networks, *Theoretical Economics*, Vol. 9: 41-97.

Acemoglu, D., Ozdaglar, A., (2011), Opinion dynamics and learning in social networks, *Dynamic Games and Applications* 1: 3-49.

Akcigit, U., Liu, Q. (2015), The role of information in innovation and competition, *Journal of the European Economic Association*, doi: 10.1111/jeea.12153.

Amir, R., Evstigneev, I., Wooders, J. (2003), Noncooperative versus cooperative R&D with endogenous spillover rates, *Games and Economic Behavior* 42: 183-207.

Arieli, I., Mueller-Frank, M., (2015), Inferring beliefs from actions, *working paper*, Israel Institute of Technology and University of Navarra.

Bergemann, D., Välimäki, J. (2008), Bandit problems, in: *The New Palgrave Dictionary of Economics*, Edited by Durlauf, S., Blume, L., Palgrave Macmillan.

Besanko, D., Wu, J., (2013), The impact of market structure and learning on the trade-off between R&D competition and cooperation, *The Journal of Industrial Economics* (LXI): 166-201.

Besanko, D., Wu, J., (2010), Government subsidies for research programs facing "if" and "when" uncertainty, *working paper*, Northwestern University and University of Arizona.

Bessen, J., Maskin, E. (2009), Sequential innovation, patents, and imitation, *RAND Journal of Economics* Vol. 40 No. 4: 611-635.

Bessen, J., Meurer, M., (2008), Patent failure: How judges, bureaucrats, and lawyers put innovators at risk, *Princeton University Press*.

Bimpikis, K., Drakopolous, (2014), Disclosing information in strategic experimentation, *working paper*, Stanford University and Massachusetts Institute of Technology.

Bolton, P., Harris, C. (1999), Strategic experimentation, *Econometrica* 67: 349-374.

Bonatti, A., Hörner, J. (2011), Collaborating, *The American Economic Review* 101: 632-663.

Bonatti, A., Hörner, J. (2015), Learning to disagree in a game of experimentation, *Cowles Foundation Discussion Paper* No. 1991. Available at SSRN: <http://ssrn.com/abstract=2573222> or <http://dx.doi.org/10.2139/ssrn.2573222>

Bramoullé, Y., Kranton, R. (2007), Public goods in networks, *Journal of Economic Theory* 135(1): 478-494.

Bramoullé, Y., Kranton, R., D'Amours, M. (2014), Strategic interaction and networks, *The American Economic Review*, 104(3): 898-930.

Calvo-Armengol, A., Jackson, M. (2004), The effects of social networks on employment and inequality, *The American Economic Review*, 94: 426-454.

Chen, Y. (2012), Innovation frequency of durable complementary goods, *Journal of Mathematical Economics* 48: 407-421.

Conley, T.G., and C. Udry, (2010), Learning about a new technology: Pineapple in Ghana, *American Economic Review* 100: 35-69.

Das, K. (2013), Strategic experimentation with heterogeneous firms and payoff externalities, *University of Exeter Economics Department Discussion Papers Series* Nr. 13/15.

Department for Business, Innovation and Skills (2010), The allocation of science and research funding: 2011/12 to 2014/15, Retrieved from www.gov.uk/government/uploads/

system/uploads/attachment_data/file/32478/10-1356-allocation-of-science-and-research-funding-2011-2015.pdf

Federal Trade Commission (2011), The evolving IP market place: Aligning patent notice and remedies with competition, Retrieved from www.ftc.gov/sites/default/files/documents/reports/evolving-ip-marketplace-aligning-patent-notice-and-remedies-competition-report-federal-trade/110307patentreport.pdf

Gale, D., Kariv, S., (2003), Bayesian learning in social networks, *Games and Economic Behavior* 45: 329-346.

Gancia, G., Zilibotti, F. (2005), Horizontal innovation in the theory of growth and development, in: *Handbook of Economic Growth*, Vol. 1A, Edited by Aghion, P., Durlauf, S.: 110-170.

Goyal, S. (2009), *Connections: An Introduction to the Economics of Networks*, Princeton University Press.

Halac, M., Kartik, N., Liu, Q. (2013), Optimal contracts for experimentation, *working paper*, Columbia University and University of Warwick.

Heidhues, P., Rady, S., Strack, P. (2015), Strategic experimentation with private payoffs, *Journal of Economic Theory* Vol. 159(A): 531-551.

Henderson, R., Cockburn, I. (1998), Scale, scope, and spillovers: The determinants of research productivity in drug discovery, *The RAND Journal of Economics* 27: 32-59

Hopenhayn, H., Squintani, F. (2016), Patent rights and innovation disclosure, *Review of Economics Studies* 83: 199-230.

Hunt, R. (2004), Patentability, industry structure, and innovation, *The Journal of Industrial Economics* Vol. 52 (3): 401-425.

Jackson, M. (2008), *Social and Economic Networks*, Princeton University Press.

Keller, G., Rady, S., Cripps, M. (2005), Strategic experimentation with exponential bandits, *Econometrica* 73: 39-68.

Keller, G., Rady, S., (2010), Strategic experimentation with Poisson bandits, *Theoretical Economics* 5: 275-311.

Lemley, M., Shapiro, C. (2005), Probabilistic patents, *Journal of Economic Perspectives* 19(2): 75-98.

Miyagiwa, K., Ohno, Y., (2002), Uncertainty, spillovers and cooperative R&D, *International Journal of Industrial Organization* Vol. 20(6): 855-876.

Mossel, E., Sly, A., Tamuz, O. (2015), Strategic learning and the topology of social networks, *Econometrica* 83(5): 1755-1794.

Nobel, R., Birkinshaw, J. (1998), Innovation in multinational corporations: Control and communication patterns in international R&D operations, *Strategic Management Journal* 19: 479-496.

Park, J. (2004), International and intersectoral R&D spillovers in the OECD and East Asian economies, *Economic Inquiry* 42 (4): 739-757.

Rocket, K., (2010), Property rights and invention, in *the Handbook of the Economics of Innovation*, Edited by Hall, B. H., Rosenberg, N., Vol. 1: 315-380.

Rosenberg, D., Solan, E., Vieille, N. (2009), Information externalities and emergence of consensus, *Games and Economic Behavior* 66: 979-994.

Sadler, E. (2015), Innovation adoption and collective experimentation, *working paper*, New York University.

Scotchmer, S., Green, J. (1990), Novelty and disclosure in patent law, *The RAND Journal of Economics* 21(1): 131-146.

Steurs, G., (1995), Inter-industry spillovers: What difference do they make?, *International Journal of Industrial Organization* Vol. 13(2): 249-276.

Teece, D. (1994) Firm organization, industrial structure, and technological innovation, *Journal of Economic Behavior & Organization* 31(2): 193-224

The National Science Foundation (2015), Uncovering Alzheimer's complex genetic networks, Retrieved from www.nsf.gov/discoveries/disc_summ.jsp?cntn_id=134049

The National Science Foundation, Cyberinfrastructure framework for 21st century Science and Engineering. Retrieved from www.nsf.gov/cif21

Wong, T. (2015), Free riding and duplication in R&D, *working paper*, The Pennsylvania State University.

Young, A. (1993), Substitution and complementarity in endogenous innovation, *The Quarterly Journal of Economics* 108 (3): 775-807.