

University of Groningen

To Complete a Puzzle, You Need to Put the Right Pieces in the Right Place

Kok, Holmer Jan

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Publisher's PDF, also known as Version of record

Publication date:
2018

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Kok, H. J. (2018). *To Complete a Puzzle, You Need to Put the Right Pieces in the Right Place: Exploring Knowledge Recombination and the Creation of New Inventions*. [Thesis fully internal (DIV), University of Groningen]. University of Groningen, SOM research school.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Chapter 5. General Discussion

The focal objective of this dissertation was to deepen our understanding of the concept of knowledge recombination. To this end, we conducted three empirical projects (chapter 2-4) in the fuel cell industry, addressing important research gaps in knowledge recombination literature. In this chapter, we first provide a short overview of the findings that emerged from the three projects. Second, we discuss our contributions to the initial research objective of this dissertation. Third, we provide a thought-provoking discussion of how our findings contribute to extant practice, paying special attention to open innovation and resource management strategies. Fourth, reflecting upon the empirical design of the three projects, we discuss the empirical contributions of this dissertation. Finally, we conclude with developing an agenda for future research on knowledge recombination.

5.1. Overview of findings

In chapter 2, we explored the influence of recombinant lag – i.e. the time that recombined components have remained unused – on the technological value of inventions. Our findings indicate that the relationship between recombinant lag and the technological value of inventions is more complex than initially anticipated. Across the two industries that we examined (i.e. fuel cell and wind energy), we found that recombinant lag has an unexpected U-shaped relationship with the technological value of inventions, signifying that low and high recombinant lag can both be associated with higher technological value. Moreover, with regards to the moderation effect of frequency of reuse on this main relationship, we found some interesting differences between the two industries. Whereas in fuel cells the U-shaped relationship between recombinant lag and the technological value of inventions mainly emerges when prior frequency of reuse is limited, in wind energy the moderation effect was not robust.

In chapter 3, we studied the impact of knowledge pool applicability – i.e. the extent to which components in the knowledge pool can be used in different application domains – on the focal firm’s partner-specific recombination within R&D alliance dyads. Holding constant the partner’s knowledge pool size, diversity,

Chapter 5

and distance, we found robust evidence that the partner's knowledge pool applicability has the expected inverted U-shaped relationship with the focal firm's partner-specific recombination. Unexpectedly, however, instead of a positive and linear relationship, a robust U-shaped relationship emerged between the focal firm's knowledge pool applicability and its partner-specific recombination.

In chapter 4, we studied the relative problem-solving performance benefits of going-together compared to going-alone in challenge-based R&D projects. In this chapter, we argued that not every organization is able to fully reap the knowledge recombination benefits of going-together. For our baseline hypothesis, we found evidence that going-together, on average, yields higher problem-solving performance than going-alone. At the same time, we found that there are certain characteristics of the focal organization that substantially influence the size of this positive problem-solving performance gap. First, in contrast to our expectations, we found that the problem-solving performance of firms is higher when going-together compared to going-alone, whereas this difference was not statistically significant for research organizations. Second, the findings indicated that the positive problem-solving performance gap between going-together and going-alone is larger for focal organizations with a greater internal knowledge pool. Finally, we unexpectedly found that a challenge-based R&D project portfolio of intermediate size is most conducive to a large positive problem-solving performance gap between going-together and going-alone.

5.2. Contributions to initial research objective

In this dissertation, our primary objective was to generate new insights into core aspects of knowledge recombination. To this end, we conducted three projects in which we ventured well beyond the traditional scope of knowledge recombination studies, generating insights that challenge many previously-held assumptions about knowledge recombination. In the following section, we provide an elaborate discussion of how this dissertation contributes to an improved understanding of knowledge recombination by generating new insights into (i) path-dependent effects that determine components' contemporary recombinant value, (ii) the

applicability of components in the knowledge pool, and (iii) combinative capabilities in interorganizational collaboration settings.

5.2.1. Path-dependent knowledge reuse trajectories

A large body of research on knowledge recombination is concerned with identifying the origins of valuable inventions (Capaldo *et al.*, 2017; Kelley *et al.*, 2013; Schoenmakers & Duysters, 2010). This research principally focuses on the original attributes of components (e.g. Phene *et al.*, 2006; Miller *et al.*, 2007; Rosenkopf & Nerkar, 2001) – i.e. attributes that are embedded into components at the time of creation. As such, they adopt a theoretical perspective in which it is implicitly assumed that the recombinant value of knowledge components is largely pre-determined at creation. In contrast to this theoretical perspective, several recent studies have argued that a component's recombinant value is not necessarily pre-determined at creation, but may actually change considerably over time. These changes, they argue, principally occur via component reuse (Belenzon, 2012; Katila & Chen, 2008; Yang *et al.*, 2010). In particular, each time a component is reused, new reuse information flows are produced which inventors can access in order to improve the subsequent recombination of this component (Katila & Chen, 2008). Therefore, a component's current recombinant value is largely a function of how it has been reused in prior knowledge recombination efforts (Yayavaram & Ahuja, 2008). In other words, there are strong path-dependent effects in knowledge recombination that determine the contemporary recombinant value of components.

Our findings in chapter 2 underline the impact of these path-dependent component reuse effects. In this chapter, we deviated from prior research, which has principally examined the frequency of reuse (i.e. the magnitude of reuse information flows), by pointing to the timing of reuse – i.e. when was a component reused during its knowledge reuse trajectory – as a crucial determinant of a component's current recombinant value. Grounded in the notion that components are recombined in different ways over time (Dosi, 1982), we developed novel theoretical arguments suggesting that precisely when an instance of reuse occurs may largely influence the contents of generated reuse information flows. Relying on organizational learning theory (e.g. Argon & Miron-Spektor, 2011), we argued

Chapter 5

that, when a component was recently reused, its recombinant potential may be rejuvenated. That is, recent reuse allows inventors to access information about how components should be contemporarily applied in knowledge recombination, resulting in inventions that fit extant technological standards. In support of this argument, our empirical tests across two industries (i.e. fuel cell and wind energy) showed that low recombinant lag (i.e. recent reuse) is associated with the creation of valuable inventions,

In addition to this rejuvenation effect, we unexpectedly found that recombining components that were last reused a long time ago (i.e. high recombinant lag) may result in inventions with considerable technological value. Performing an extensive post-hoc exploratory analysis, we explained this alternative finding by arguing that reuse information flows attached to temporally distant instances of reuse may contain information about the component that was simply too far ahead of its time (Garud & Nayyar, 1994), subsequently becoming ‘frozen’ for prolonged periods. Decades later, however, during the emergence of a new technological cycle within an industry (Tushman & Anderson, 1986), inventors may be able to ‘defrost’ these reuse information flows, attach new meanings to them (Sonenshein, 2014), and subsequently leverage them to generate inventions with considerable technological value.

In sum, we emphasize that, in order to obtain a fuller understanding of a component’s current recombinant value, it is not sufficient to only consider the attributes that are embedded into the component at creation. Instead, our findings in chapter 2 point out that the current recombinant value of a component is, to a large extent, determined by the path-dependent knowledge reuse trajectory through which it travelled from inception until the present. More specifically, we contribute to extant literature by emphasizing that even when two components were created at the same time, they may go through considerably different knowledge reuse trajectories over time, where not only the number of times they were reused (i.e. frequency of reuse) matters, but also the time that they have remained unused (i.e. recombinant lag). We therefore encourage future studies on knowledge recombination to account for the temporal dimension of component reuse, integrating recombinant lag into their conceptual and empirical framework.

5.2.2. Applicability of components in the knowledge pool

R&D scholars extensively examine the relation between partners' knowledge pool characteristics and the focal firm's internal knowledge recombination activities (e.g. Nooteboom *et al.*, 2007; Sampson, 2007; Phelps, 2010). Both conceptually and empirically, these studies tend to focus on aggregate characteristics of the partner's knowledge pool, such as its diversity or size. Hence, these studies analyze partners' knowledge pool holistically, rather than to isolate individual components that constitute the knowledge pool. Adopting this research approach, numerous alliance studies (e.g. Schilling & Phelps, 2007; Phelps, 2010; Wuyts & Dutta, 2014) assume that components situated within large and highly technologically diverse knowledge pools are more broadly applicable, creating various opportunities to generate new combinations (Ahuja & Katila, 2001; Fleming, 2002). However, findings have been rather inconsistent in this research area: whereas some studies detect a strictly positive influence of partners' technological diversity on the focal firm's knowledge recombination activities (e.g. Phelps, 2010; Subramanian & Soh, 2017), others find inverted U-shaped (e.g. Vasudeva & Anand, 2011) and U-shaped effects (e.g. Wuyts & Dutta, 2014). It therefore seems that our understanding of the knowledge recombination implications of component applicability within R&D alliances is currently incomplete.

In chapter 3, we made important progress towards resolving these extant inconsistencies. In particular, relying on recent knowledge recombination insights (e.g. Dibiaggio *et al.*, 2014; Wang *et al.*, 2014), we argued that individual components that constitute the knowledge pool may actually differ substantially in terms of applicability. In this chapter, we therefore shifted the theoretical lens from aggregate knowledge pool characteristics to individual components within the knowledge pool. More specifically, we focused on the knowledge pool applicability of R&D alliance partners – i.e. the extent to which individual components situated in the knowledge pool can be used in different application domains of a technological field. From the focal firm's perspective, we developed novel theoretical arguments regarding the knowledge recombination benefits and liabilities of collaborating with partners with higher knowledge pool applicability. We theorized an inverted U-shaped relationship between the partner's knowledge pool applicability and the focal firm's intensity of partner-specific recombination.

Chapter 5

Here, the upward slope of the relationship is driven by the additional flexibility that the focal firm gain in terms of where and how it may apply components accessed from partners in knowledge recombination (Dibiaggio *et al.*, 2014; Wang *et al.*, 2014; Yayavaram & Ahuja, 2008), and the downward part of the slope mainly arises from learning complexities associated with overly-applicable component knowledge (Fleming & Sorenson, 2001; Hargadon & Sutton, 1997). In our empirical analyses, holding constant the size and diversity of the partner's knowledge pool, we found strong support for this hypothesis: a robust inverted U-shaped relationship between the partner's knowledge pool applicability and the focal firm's intensity of partner-specific knowledge recombination consistently emerged.

Taken together, chapter 3 demonstrates the importance of accounting for variance in the applicability of individual components within the knowledge pool of R&D alliance partners. Different from extant alliance research, which infers components' level of applicability from the size and technological diversity of partners' entire knowledge pool (e.g. Lahiri & Narayanan, 2013; Phelps, 2010; Schilling & Phelps, 2007), we make a strong case for digging deeper into partners' knowledge pool, examining whether individual components are, in fact, broadly applicable or not. Adopting a similar research approach, future research may be able to resolve important inconsistencies regarding the relationship between partners' knowledge pool diversity and the focal firm's internal knowledge recombination activities. Besides this, on a more general note, our study showed the value of adopting state-of-the-art knowledge recombination insights when examining the inventive benefits of R&D alliances, a research approach which is only rarely used in extant alliance research.

5.2.3. Combinative capabilities in interorganizational collaboration

Knowledge recombination studies often implicitly assume that components, to which the focal organization gains access via interorganizational collaboration, are directly connected to its inventive output. At the same time, it has been shown that, even when two organizations have access to the same set of components, there is typically much variation in terms of the quantity and quality of new combinations that they produce (Dibiaggio *et al.*, 2014; Yayavaram & Ahuja, 2008). These

variations in inventive output arise from organizations' idiosyncratic abilities to identify valuable component knowledge in partners' knowledge pool, retrieve it by developing a solid understanding of its recombination characteristics, and subsequently use it in new recombination efforts to create new inventions (Fleming & Sorenson, 2001; Zahra & George, 2002). In this dissertation, we developed important new insights about these combinative capabilities (Kogut & Zander, 1992), examining how organizations leverage them in order to enhance the inventive benefits of interorganizational collaboration.

In chapter 3, we argued that focal firms with higher knowledge pool applicability, having prior experience creating component knowledge with broad applicability, can more intensively recombine the partner's components in R&D alliances. We argued that this experience can help firms overcome two barriers commonly blocking the recombination of alliance partners' component knowledge: (i) perceived exhaustion of recombinant potential and (ii) the pursuit of ultimately fruitless recombination opportunities. Tackling the former barrier, we argued that focal firms with higher knowledge pool applicability have, in general, a better understanding of how components can be recombined in different ways (Boh *et al.*, 2014). As such, these types of firms are able to elevate the recombinant potential of components in the partner's knowledge pool, envisioning a greater number of ways in which these components can be used in knowledge recombination (Boh *et al.*, 2014; Fleming, 2002; Henderson, 1995; Yang *et al.*, 2010). Moreover, tackling the latter barrier, we argued that experience with building widely-applicable component knowledge provides the focal firm with a better understanding of where the recombinant limits of components lie (Henderson, 1995; Yang *et al.*, 2010). Hence, while pursuing recombination opportunities in the partner's knowledge pool, these focal firms are able to quickly locate valuable components that provide opportunities to generate new combinations and not waste time on components which are "technological dead-ends" (Podolny & Stuart, 1995: 1225), bypassing time-costly trial-and-error activities in the process (Fleming, 2001; Katila & Ahuja, 2002). In line with these theoretical arguments, our findings showed that, beyond a certain threshold value, the focal firm's knowledge pool applicability is indeed positively associated with its intensity of partner-specific recombination.

Chapter 5

In chapter 4, we further probed the usefulness of combinative capabilities, and compared the problem-solving performance of focal organizations that go-together or go-alone in challenge-based R&D projects. Prior studies on challenge-based R&D projects, following the open innovation paradigm (Chesbrough, 2006), often implicitly assume that going-together always outperforms going-alone in terms of generating high-quality technological solutions (e.g. Olsen *et al.*, 2016). These studies argue that, by going-together, focal organizations gain access to a potentially larger and more diverse set of components, allowing to create higher quality technological solutions (Das & Teng, 2000; Olsen *et al.*, 2016). In chapter 4, however, we pointed out that the extent to which these potential recombination opportunities can actually be realized depends on the ability of the focal organization to identify, retrieve, and recombine partners' components (Wuyts & Dutta, 2014; Zahra & George, 2002). Our findings strongly support this contention, and cast significant doubts over the purported advantages of going-together strategies to address extant societal challenges, indicating that, for many focal organizations, going-alone strategies often perform equally well in terms of problem-solving performance.

Jointly, in this dissertation we enriched our understanding of the value of combinative capabilities in interorganizational collaboration settings. We emphasized that, rather than to assume a direct connection between potential knowledge inputs accessed from partners and the focal organization's own inventive outputs, it is important to take into account that not all organizations are equally capable of leveraging the available set of components for knowledge recombination (Wuyts & Dutta, 2014; Zahra & George, 2002). In chapter 3, we conceptually focused on two known impediments to knowledge recombination, developing new theoretical arguments about how focal firms may develop abilities to successfully recombine the component knowledge of R&D alliance partners. In chapter 4, we found that strong combinative capabilities play a crucial role in translating potential into realized knowledge recombination opportunities when going-together in challenge-based R&D projects. The findings from chapter 4 are striking, as they imply that interorganizational collaboration, at least within the scope of challenge-based R&D projects, may not be a worthwhile strategy to pursue for focal organizations lacking the right combinative capabilities. For these types

of focal organizations (most notably, research organizations and focal organizations with small internal knowledge pools), a going-alone strategy may be equally-beneficial when it comes to solving extant technological problems within a field.

5.3. Practical contributions

Next to theoretical contributions, the three empirical projects on knowledge recombination generated important practical contributions. These practical contributions fall into two categories, which we discuss in the following section: (i) open innovation strategies for new technology production and (ii) resource management strategies for knowledge recombination.

5.3.1. Open innovation strategies for new technology production

In line with the emerging open innovation paradigm (Chesbrough, 2006), academics and practitioners have called for more attention to collaboration across organizational boundaries, stimulating the exchange of heterogeneous resources and capabilities between different organizations (e.g. firms, universities, government institutions). We contribute to this emerging paradigm in important ways. In chapter 2, we found that available reuse information flows substantially influence the value of subsequent inventions. Especially when reuse information flows were recently produced, inventors become able to recombine associated components more effectively, leading to more valuable inventions as a result. Hence, we encourage policy-makers to push for more transparency in the production of new inventions, such that reuse information flows become more easily available to inventors. In this way, there may also be higher societal returns to investments into R&D, since the value-added benefits of new inventions do not singly remain within the confines of originating organizations (Ahuja, Lampert, & Novelli, 2013; Murray & O'Mahony, 2007; Yang *et al.*, 2010).

In chapters 3 and 4, we focused on how organizations use interorganizational collaboration activities to enhance their own knowledge recombination activities. Based on the extant literature (e.g. Gomes-Casseres *et al.*, 2006; Grant & Baden-Fuller, 2004; Mowery *et al.*, 1996), we argued that focal organizations often engage in interorganizational collaboration with the intention

Chapter 5

to access the component knowledge of other organizations. However, especially in the chapter 4, we found strong evidence to suggest that these benefits may not materialize for certain organizations. Hence, we encourage practitioners to carefully consider which organizations may actually be able to capture the knowledge recombination benefits of interorganizational collaboration. Especially when designing programs addressing grand societal challenges, a more careful consideration of organizations' ability to benefit from going-together strategies seems warranted.

5.3.2. Managing resources for knowledge recombination

In this dissertation, we consistently found that available resources, within and between organizations, should be carefully managed such that available knowledge recombination opportunities can be fully realized by the focal organization. In chapter 2, strong indications emerged that much technological value may be realized from recombining dormant components and accessing their associated reuse information flows. Therefore, we argue that the existing knowledge stock should continuously be reevaluated (Garud & Nayyar, 1994), such that these 'hidden gems' may be timely identified. For practitioners, this implies that they should more fully consider how previously generated resources may be managed such that they could have a long-lasting impact on organizations' ability to generate new value.

In chapter 3, we found, unexpectedly, that capabilities emerging from developing widely-applicable component knowledge initially decrease the focal firm's intensity of partner-specific recombination. In chapter 4, similar findings emerged, as we found that the internal knowledge pool size of the focal organization, initially, does not substantially increase the relative problem-solving performance benefits of going-together compared to going-alone. We explained these alternative findings by arguing that organizations need to 'learn to learn' (Levinthal & March, 1993); that is, they need to better understand how prior knowledge recombination outputs translate into capabilities that improve future knowledge recombination efforts (Ghosh *et al.*, 2014; Lewin *et al.*, 2011; Wuyts & Dutta, 2014). Thus, practitioners should carefully manage prior knowledge recombination activities and their associated inventive outputs (Anand *et al.*, 2016;

Ghosh *et al.*, 2014; Marsh & Stock, 2006), carefully considering how these prior experiences can be leveraged to improve knowledge recombination activities in interorganizational collaboration settings (Subramanian & Soh, 2017).

5.4. Empirical contributions

In this dissertation, we used state-of-the-art empirical methods to analyse rich quantitative data on the inventive activities of organizations in the fuel cell industry. In each project, we also introduced several innovations in the design of the empirical approach, which we discuss in the following section.

In chapter 2, we developed two major empirical contributions. First, following recent patent studies (e.g. Bakker *et al.*, 2016; de Rassenfosse *et al.*, 2013; Nakamura *et al.*, 2015), we collected patent applications related to fuel cell technology from every patent office in the world, and subsequently aggregated these single patent office applications to the patent family level (using the EPO's DOCDB definition). Studying patent families was particularly important in the fuel cell industry, as numerous prominent players in this industry are non-American firms that tend to patent extensively in their home region's patent office (e.g. Toyota at the JPO and Renault at the EPO). Moreover, we relied on patent citations to track knowledge recombination and reuse (Katila & Chen 2008; Yang *et al.*, 2010), and studies show that patent families provide a better account of patent citations than single patent office applications (Albrecht *et al.*, 2010; Bakker *et al.*, 2016; Nakamura *et al.*, 2015). In this way, we deviated from prior patent-based studies on knowledge recombination, which singly rely on patent applications from either the USPTO (e.g. Fleming, 2001; Miller *et al.*, 2007; Nemet & Johnson, 2012) or the EPO (e.g. Gruber *et al.*, 2013; Schoenmakers & Duysters, 2010). Second, in chapter 2, we had some concerns regarding the generalizability of our main findings from the fuel cell industry. Therefore, we collected additional data from a completely different industry (i.e. wind energy), in which we were able to confirm most, but not all, of our findings. In this way, we were able to further substantiate our contributions to the knowledge recombination literature (Bettis, Helfat, & Shaver, 2016).

Chapter 5

In chapter 3, we made empirical contributions to studies on R&D alliances and patenting activities in general. First, we described how R&D alliance activities are heavily underestimated when relying on popular alliance databases, such as Thompson Reuters' SDC Platinum Joint Venture and Strategic Alliances database (Lavie, 2007; Lavie & Rosenkopf, 2006). To address this issue, we collected data on all fuel cell R&D alliances formed between 1978 and 2007 by examining close to 50,000 news articles within the LexisNexis database, following a handful of prior studies that have adopted a similar approach (e.g. Ahuja, 2000; Phelps, 2010; Vasudeva & Anand, 2011). Second, we tracked R&D alliances over time (Ahuja, 2000; Hashai *et al.*, 2018; Phelps, 2010), in order to estimate their starting and ending dates more precisely. This approach deviates from more conventional methods employed by alliance studies, in which a fixed lifespan (ranging from 2 to 5 years) is usually assumed for every alliance (e.g. Schilling & Phelps, 2007; Vasudeva & Anand 2011). Our approach is especially useful when we consider that there is actually much heterogeneity in the lifespans of alliances (Deeds & Rothaermel, 2003). For example, in our sample, we detected that some R&D alliances only lasted for one year, whereas others lasted for more than 15 years. Third, we conducted an in-depth examination of fuel cell patents' IPC codes to derive information about the applicability of knowledge components. Normally, patent IPC codes are aggregated at the firm-level, with little regard for the actual content of these codes. In this chapter, instead, we inspected several leading fuel cell review articles and news articles, and examined numerous patent documents, in order to identify the relevant application domains of fuel cell technology and their associated patent IPC codes.

Finally, in chapter 4, we used data from the U.S. Department of Energy's Hydrogen and Fuel Cells Program to examine the difference in problem-solving performance between focal organizations that go-together or go-alone in challenge-based R&D projects. Using this data, we developed several important empirical contributions. First, we adopted a performance metric that is (i) directly connected to the outputs of the project, (ii) provided by peer reviewers that are objective, anonymous, and experts in their respective fields, and (iii) comparable across projects. This performance indicator is arguably superior to the ones used by prior studies in which, for instance, firm-level performance data (e.g. net profit

margin, stock market returns) is used to infer something about the outputs of particular collaborative projects (e.g. Jiang *et al.*, 2010; Sampson, 2007). Second, the challenge-based R&D projects that we examined, in which focal organizations go-alone or go-together, are highly comparable with one another, allowing us to tease out the problem-solving benefits of going-together strategies relative to going-alone strategies. Moreover, in our sample, we observed that these two R&D approaches are employed almost equally often by project leaders. Finally, we measured the focal organization's internal knowledge pool size, not by aggregating every single patent that the focal organization filed in the past, but instead by focusing on granted patents that were filed within the scope of the HFPC. In this way, our measure more accurately captured patents that were actually relevant to the particular challenge-based R&D projects that we studied.

5.5. Future research directions

In this section, we discuss several promising avenues for future research on knowledge recombination that emerged from our results in chapters 2-4. First, it is often argued that organizations need to use collaborative strategies in order to change the size, composition, and configuration of their own knowledge pool (Rosenkopf & Almeida, 2003). At the same time, in chapters 3 and 4, we already found strong indications that not every organization is actually able to access and recombine partners' external components very effectively. Future studies should build upon these findings, and examine to what extent organizations might be better off solely relying on internally available component knowledge to generate new inventions (Baker & Nelson, 2005; Miller *et al.*, 2007). As argued by recent knowledge recombination studies (e.g. Dibiaggio *et al.*, 2014; Yayavaram & Ahuja, 2008), some organizations are able to generate a considerable number of new inventions by reconfiguring existing component combinations within their own knowledge pool (Henderson & Clark, 1990). It would be interesting to compare the relative benefits of such inward-oriented strategies, with more outward-oriented strategies in which direct collaborations with other organizations are considered.

Second, in this dissertation, we consistently developed our theoretical arguments from the perspective of one focal entity (i.e. the focal inventor in chapter

Chapter 5

2, the focal firm in chapter 3, and the focal organization in chapter 4). However, in many cases, such as when organizations collaborate in knowledge recombination, new combinations represent a recombination of different organizations' individual components. Hence, to enrich our understanding of how collaborative activities should be optimally configured, future research should study precisely whose components are actually utilized in collaborative knowledge recombination activities.

Third, future studies should examine which factors allow organizations to recombine components with specific attributes. We made some important progress in this area already, showing in chapter 3 that focal firms have different abilities that allow them to recombine alliance partners' knowledge components more intensively. Nevertheless, the body of research that focuses on these antecedents of knowledge recombination is still rather small (Carnabuci & Operti, 2013; Hohberger, 2014; Phelps, 2010), implying that there are still many 'low-hanging fruits' waiting to be picked in this research area. For example, an obvious extension of our project on recombinant lag would entail examining which factors stimulate organizations' ability to recombine components with different levels of recombinant lag and frequencies of reuse.

Finally, on the empirical front, it would be interesting to delve more deeply into non-patent-based measures of knowledge recombination. In our research, we relied extensively on patent citations (chapters 2 and 3) to track knowledge recombination. We recognized the many advantages of this method (e.g. the ability to assign components to particular originating firms, the ability to easily track knowledge recombination over time), but also acknowledged some limitations (e.g. bias created from examiner-added citations) (Jaffe & de Rassenfosse, 2017). Some qualitative studies on knowledge recombination already exist (e.g. Hargadon & Sutton, 1997; Fleming, 2002; Sonenshein, 2014), but large-scale quantitative research on knowledge recombination that does not principally rely on patent- or publication-based data is still largely missing (an exception is: Sidhu, Commandeur, & Volberda, 2007). At the very least, such quantitative research efforts would allow us to verify the generalizability of existing studies that use patent-based measures, of which there are now many (Laursen, 2012; Sidhu *et al.*, 2007).

5.6. Concluding thoughts

In this dissertation, we showed that the creation of new inventions is much like solving a puzzle: you need to put the right puzzle pieces (components) in the right place (combinative capabilities) in order to complete the puzzle (invention). In chapter 2, we showed that the recency of component reuse influences whether inventors are able to generate technological value from the recombination of particular components. This underlines the fact that knowledge recombination is a path-dependent process, in which a component's current usefulness in knowledge recombination cannot be understood separately from its historical knowledge reuse trajectory. Moreover, in chapter 3, we showed that components are highly malleable, with different application domains attached to them. Within the R&D alliance context, differences in component applicability may significantly impact opportunities of firms to engage in knowledge recombination. Finally, in chapter 4, we provided much-needed nuance to ongoing discussions regarding the benefits of collaborative knowledge recombination strategies. We found that organizations require adequate combinative capabilities in order to reap the knowledge recombination benefits of going-together. We hope that these findings provide interesting starting points for future academic research on knowledge recombination. Moreover, we are hopeful that our contributions to practice inspire practitioners and policy-makers alike to more carefully consider the pros and cons of various open innovation strategies, and the management of resources for new invention production.

