



Paper to be presented at the DRUID 2012

on

June 19 to June 21

at

CBS, Copenhagen, Denmark,

Learning to Write Killer Apps? A Specialist/Trans-Specialist Perspective in Developing Innovations for the Marketplace

Terence P C Fan

Singapore Management University
LKC School of Business
terencefan@smu.edu.sg

Xuesong Geng

Singapore Management University
LKC School of Business
xsgeng@smu.edu.sg

Abstract

Commercially successful but not necessarily high-tech innovations can change the fortunes of firms and the quality of lives of many. This paper argues that while groups of specialists can quickly bring together knowledge from multiple domains in developing and implementing innovative ideas, a typical lack of system-wide perspective would prove to be a constraint over successive innovations. In contrast, a single individual who acquires specialist knowledge in multiple domains would avoid this constraint, but would incur a penalty to acquire knowledge across domain boundaries upfront. This leads to two opposing performance predictions on the performance of successive innovations: an increasing trend at a decreasing pace for the former and a J-shape for the latter. Moreover, experience developing solo innovations strengthens the performance for subsequent group innovations. These hypotheses are supported by empirical data on the download performance of third-party software applications written for use within the Facebook social networking context.

Jelcodes:O31,-

Learning to Write Killer Apps?

A Specialist/Trans-Specialist Perspective in Developing Innovations for the Marketplace

Abstract : Commercially successful but not necessarily high-tech innovations can change the fortunes of firms and the quality of lives of many. This paper argues that while groups of specialists can quickly bring together knowledge from multiple domains in developing and implementing innovative ideas, a typical lack of system-wide perspective would prove to be a constraint over successive innovations. In contrast, a single individual who acquires specialist knowledge in multiple domains would avoid this constraint, but would incur a penalty to acquire knowledge across domain boundaries upfront. This leads to two opposing performance predictions on the performance of successive innovations: an increasing trend at a decreasing pace for the former and a J-shape for the latter. Moreover, experience developing solo innovations strengthens the performance for subsequent group innovations. These hypotheses are supported by empirical data on the download performance of third-party software applications written for use within the Facebook social networking context.

Introduction

In the increasingly innovation-driven world, creative ideas that, upon implementation, turn out to be commercially successful can change the fortunes of firms and the quality of lives for many (Amabile, 1996). While few successful innovations – either scientific or entrepreneurial in nature – have been developed by individuals or organizations as their first and one-off attempts (Parker 2006; Yayavaram & Ahuja, 2008), the extent to which past innovations contribute to the commercial success of future ones remains an important area of inquiry. Developing innovations that lead to commercial success requires knowledge in multiple disciplines – at least the technology that underlies the end-product and the knowledge of how the end-product may be adopted by others. Since individuals do not always lend themselves to become specialists in multiple disciplines, specialists of different knowledge domains often collaborate in groups – each working to advance in one’s own knowledge domain – to enhance the commercial performance of the end-product. The pattern of how experience impacts the commercial performance of future innovations may therefore be different for individuals working as specialists in a group versus those who must develop expertise in multiple disciplines as they work alone in their innovation projects. This is the subject of investigation in this paper.

Scholars studying the development of scientific patents report how knowledge development takes place competitively at the state-of-the-art knowledge frontiers (e.g., Audia & Goncalo, 2007; Singh & Fleming, 2010), but a vast number of commercially successive ideas, including comic books and country music, rely on far less advanced technologies (Taylor & Greve 2006; Lingo & O’Mahony 2010). We focus on this latter category of innovations.

While scholars have not tackled the outlined issues outright, many have examined different facets of the problem – underscoring how experience matters in innovation

performance. In the search for new ideas, recombining or refining pre-existing ideas have been shown to be effective in developing scientific patents (Fleming 2000). Prior studies have confirmed the path-dependent nature of knowledge-intensive research, recognizing the role of experience in innovation activities (Penrose 1959; Cohen & Levinthal 1990; Perry-Smith, 2006; Fleming, Mingo & Chen 2007). However, experience in innovation-related activities does not squarely lend themselves to the traditional learning curve (e.g., Darr, Argote & Epple, 1995; Kim & Miner, 2007) that predicts future improvements at both the individual and organization levels from increased experience in the same context. This non-conformance may be attributed to the fact that by definition one innovation must differ from another in some aspects. In addition to learning to perfect techniques of implementing a specific innovative idea, innovations carry with them an implicit search as to what a new idea should be about. This added element of creativity contrasts with the predetermined task environment that typically associates with the traditional learning curve, and opens the possibility of departures from it. In particular, we believe that there are different learning dynamics for individuals with increasing specialized knowledge in one domain versus others with increasing knowledge over multiple knowledge domains. Our empirical data from third-party Facebook software applications support our hypothesis.

Innovation and Specialization

Innovations involve both the identification of a creative idea and its implementation. Based on a comprehensive literature review, Crossan and Apaydin (2010: 1155) broadly define innovations as ‘production or adoption, assimilation, and exploitation of a value-added novelty in economic and social spheres; renewal and enlargement of products, services, and markets; development of new methods of production; and establishment of new management systems.’

While Crossan and Apaydin (2010) consider innovations as both processes and outcomes, we in this study examine innovations through their outcomes in a product or service which in turn is subject to economic forces of the respective community.

Inherently, creativity is difficult to define, and management scholars have generally labelled creativity as ‘a search for novelty’ (Lampel, Lant & Shamsie, 2000: 266), and creative ideas as those that are novel and appropriate to a specific context (Amabile, 1996). The generation of new ideas is achieved invariably through combining elements of pre-existing ideas – either from the knowledge of the focal individual, or through social ties with others (Csikszentmihayi, 1988; Simonton, 1999; Perry-Smith & Coff, 2011). Studies in social networks confirm how the diversity in specialist knowledge – beyond differences in demography (Taylor & Greve 2006) – contributes to the generation of novel ideas (Cummings, 2004; Perry-Smith, 2006; Sosa, 2011). What is silent in these studies is how performance unfolds over successive innovations, or in other words, how past experience in innovations contributes to the success of future ones. We address this point by first focusing on an important assumption in these studies.

A fundamental assumption implicit in studies on innovations is the existence of specialist knowledge – beyond an elementary understanding – within a myriad of knowledge domains. Elementary knowledge can be acquired by taking an introductory class or reading relevant books. It typically takes some considerable effort to kick start one’s elementary knowledge in a new domain as its ‘ontology, concepts, and operations’ may be foreign to that individual (Postrel 2002: 319), and contextual cues may be involved in the learning process (Zander & Kogut 1995). However, once this elementary stage is completed, acquiring further specialist knowledge is relatively straightforward – accomplished by reading relevant articles in research journals or conference proceedings, by reading trade journals or befriending industry professionals, or by

taking up relevant employment – until one nears the knowledge frontier. At the knowledge frontier, acquiring more specialist knowledge becomes costly because original research is often necessary, and the result of state-of-the-art research conducted by others may not be publicly accessible due to their commercial sensitivity. Figure 1 is a stylized presentation of the economics of knowledge acquisition in a particular domain, and forms the base assumption of this paper. In particular, we focus on the downward-sloping stretch of this curve.

As a direct consequence, it is common for individuals to develop into specialists within a knowledge domain. The same tendency is mirrored in the departmentalized nature of large organizations, or the division of specialized labor in multi-disciplinary project groups – in which individuals take on the role of specialists as each work to complete part of the project.

Orthogonal to the development as a specialist in one knowledge domain is the possibility that an individual acquires only elementary knowledge but does so in multiple knowledge domains. The reason why an individual typically does not simultaneously have deep specialist knowledge in multiple domains is because of the diseconomies in learning across different knowledge domains compared to within one domain. These individuals can understand stylized specialist concepts from different knowledge domains, and can therefore act as knowledge brokers across disparate domains (Lingo & O'Mahony 2010) – as most new ideas are generated from some combination of pre-existing ideas, technologies, strategies or processes (Simonton 1988). Postrel (2002: 306) refers to these individuals as having trans-specialist capability. Henderson (1994), and Iansiti and Clark (1994) refer to this capability when manifested at the organizational level as integrative capability.

While pure 'trans-specialists' or 'integrative' individuals can act as knowledge brokers across different domains, their ability to implement any innovative ideas is thwarted by the lack

of specialist knowledge in any particular domain. This is evidenced by the fact that no large organization is known to host a large department of generalists who simply visit other departments in search of broker-able knowledge. Likewise, knowledge brokers in music production ordinarily work in the capacity of producers – with a set of detailed roles in addition to the knowledge brokering function (Lingo & O’Mahony 2010). When we refer to trans-specialists in this paper, we therefore refer to individuals with elementary knowledge in multiple knowledge disciplines but also at least a small amount of specialist knowledge in one discipline for them to be functionally capable of developing and implementing an innovative idea alone.

Specialist Innovation Experience and Performance

If the commercial performance of an (end-product of) innovation is an intended objective, the knowledge in at least three disciplines is likely needed to bring about commercial success: i) the knowledge of how the creative process works: how to initiate it, iteratively refine ideas and prudently short-list potential ideas for further implementation (Perry-Smith & Coff, 2011, found how the generation and selection of best ideas may entail specialist knowledge not normally found within a single individual, further supporting the multi-disciplinary involvement of developing innovations), ii) the technical knowledge of how to implement an idea, and iii) how certain features of an innovation may encourage others to adopt (or impede others from adopting) it. As a result of the multi-disciplinary knowledge required in developing successful innovation, multiple specialists – each specializing in different knowledge domains but mutually relatively ignorant of each other’s domain – often work together to deliver the end-product.

As soon as there is some delegation of responsibility for individuals involved in a collaborative innovation project – which is common for any group projects – the individuals

typically take on the role of specialists in the respective domains and concentrate on the assigned domain. After concluding which product idea to realize, the design objective in the respective domain of a specialist in this collaborative group is likely negotiated jointly with other such specialists but remain the core responsibility of this assigned individual. Through repeatedly working to creatively conceive and implement innovative ideas, a specialist therefore expands his/her knowledge in the specialized domain in a similar manner as in a traditional learning curve (with a reasonably defined objective to work on, etc.), leading to efficiency gains that help speed up the process of realizing an innovative idea. This manner of performance improvement is therefore expected to be similar to what is typically described as a learning curve originally derived in multiple manufacturing and service contexts (Argote & Epple 1990; Darr, Argote & Epple 1995; Haunschild & Sullivan, 2002): large performance can be achieved early on in successive projects, but marginal improvements decline over successive projects.

As an individual expands his or her specialist knowledge in a domain, he or she also develops greater specialist capability and in so doing, enables more ideas to be implemented and expands the possibility frontier of developing innovations (Postrel, 2002). Incrementally, however, increasing specialization would after some point yield declining marginal improvement, after ‘easy’ opportunities for large performance gains have been exploited – consistent with the learning curve literature. There is another, deeper reason why the performance may actually decline with deep specialization at the sub-problem (or sub-system) level. Note that as specialist knowledge increases, the specialists in different domains remain largely ignorant of the details one another’s work across domain boundaries. As each specialist improves an isolated part of the overall problem, the communication and coordination cost of the intricacies of a solution developed with deeply specialized knowledge increases across domain

boundaries, and as the performance of one subsystem increases, other sub-systems may unexpectedly become a performance constraint. For instance, the design of micro-processors has for many years involved only specialists in electrical and material engineers – working on increasingly smaller transistors and how they are arranged and etched onto circuit boards. However, the increasingly miniaturized components soon meant that too much heat was generated (not a primary concern to specialists before) when these components are in operation, thus jeopardizing the operation of the entire circuit board (Moore 2006).

Repeatedly working with the same individuals also improves performance as there is an increased understanding on ‘who knows what’ (e.g., Reagans, Argote & Brooks 2005). The effect of adopting specialist orientation that we theorize here differs from the performance effect of team familiarity. In studies of performance effects of group familiarity, organization scientists typically re-set the ‘clock’ of team familiarity whenever new collaboration partners begin working in the same project (e.g., Taylor & Greve 2006). In our perspective, as long as an individual continues to play the role of a specialist in the same knowledge domain in a group setting – whether or not the focal individual works with the same collaboration partners over time, his/her performance is impacted as a result of the increased specialization.

Certainly, involving more than one individual in projects with potential commercial rewards always carries an appropriation hazard – there would not be an exception in a group of specialists (Hsieh, Nickerson & Zenger 2007). This means that while the sum of all potential specialist knowledge among constituent individuals in a group may increase over time, the specialist knowledge contributed by an individual to that group project may represent a small set of his or her complete range of capability. In effect, the range of specialist knowledge presented by a group project with appropriation hazards may therefore represent the vision and specialist of

a lead-individual of that group plus only a small subset of others in the group, rather than the sum of the entire gamut of knowledge acquired by all constituent individuals.

This preceding discussion leads us to believe that increased experience of a specialist in a knowledge domain relevant in a subsystem design should yield an improving commercial performance for the overall system, but the rate of improvement eventually slows. This leads to the following hypothesis:

Hypothesis 1: The commercial performance of successive innovations developed by an individual acting as a specialist improves at a decreasing rate.

Trans-specialist Knowledge and Performance Implications

The potentially declining performance as multiple specialists (in different disciplines) each advances toward the respective frontiers of their knowledge domains can be traced to the increased coordination cost and complexities across domain boundaries. Each one does not know much about the other knowledge domains, and relies mostly on other specialists to communicate any critical concerns or constraints.

Senior managers in multi-business corporations deciding how to allocate capital across divisions in fact face a similar problem as our group of mutually ignorant specialists. No manager can claim to have an intimate knowledge of all of the corporation's business units, and with that the growth prospects and potential competitive and technological vulnerabilities, yet senior managers have to allocate capital to maximize profit of the entire corporation. They therefore rely on divisional managers to communicate the relevant information and forecasts – which may be strewn with the managers' own biases and difficult to independently verify. It is therefore no surprise at all that the senior managers may in aggregate and individually make their

capital allocation decisions not necessarily in the best interest of the entire corporation. In a recent study of over six hundred multi-business corporations, senior managers are shown to do exactly that – their internal capital allocations are biased toward equality over the number of business units into which the firm was partitioned (Bardolet, Fox & Lovallo 2011).

In theory at least, this essentially communication or knowledge ‘internalization’ problem can be easily solved: put each senior manager through each division and learn all there is to learn from the specialists, and then – only then – they would be able to make the best allocation decision for the firm, barring unexpected surprises from the business environment. In practice, however, this solution would take a very long time and hence not practical.

Nevertheless, the notion that having at least one individual to gain a system-wide perspective – knowing all of what the specialists know – would likely make better decisions for the end-product (or entire system) has an important implication for innovation development. Empirical evidence suggests that specialists in similar knowledge domains working on different aspects of the same problem tend to know what concerns each other might face, and therefore collectively perform better than other such groups of specialists – but their end-product tend to be sub-system components to larger systems (Reagans & McEvily 2003; Obstfeld 2005). Extrapolating this point to the development (and implementation) of innovative ideas for the marketplace, we can see how a single individual with specialist knowledge in multiple disciplines is still likely able to generate performance improvements on the end-product (at the system level) when collaborating specialists begin to see declining performance – assuming that the ‘trans-specialist’ has the right amount of specialist knowledge in the required disciplines.

Our important assumption that one trans-specialist can perform better than mutually ignorant specialists rests on his/her ability to acquire (and appropriately apply) knowledge

required from the appropriate disciplines. From our earlier example of the increasingly miniaturized components of a circuit board shows that the appropriate disciplinary knowledge to solve the sub-system level 'over-specialization' may not be known ahead of times. In the context of innovation development, where the next project is never an exact replica of the previous one, the kind of knowledge required to improve commercial performance in one project may be different from the kind of knowledge required in the next project. Learning what new knowledge (and in what domain this knowledge resides) is required to enhance the commercial performance of the next innovation, as well as what past knowledge is appropriate to apply in the next innovation therefore creates another layer of complexity and potentially false causal associations for trans-specialists working alone in an innovation project. Specialists with deep knowledge in a particular knowledge domain may alleviate this complexity because by definition deep specialist knowledge increases the range of possibility of ideas that can be materialized, and this helps compensate for the lack of clear causal mechanisms (to enhance commercial performance) encountered by trans-specialists with only shallow specialist knowledge in a particular domain.

Since knowledge in multiple disciplines are required to achieve commercial success for innovations (but more knowledge does not guarantee better performance), individuals responsible for developing these alone must by definition be capable trans-specialists with at least some specialist knowledge in a specific domain. The upfront learning cost that a trans-specialist must incur before his/her innovations achieve sustained improvement in commercial success is consistent with the observation of long time intervals between breakthrough inventions by serial innovators who work alone. For example, the inventor of Segway human transporter spent a long time developing the appropriate specialist knowledge between his successive

(commercially successful) innovations (Harris 2010). Once this learning cost is overcome (and that the correct causal link between performance and knowledge is understood), trans-specialists hold the promise of higher-quality innovations than collaborative specialists. Within at least one domain of U.S. patents – tennis racquets – habitually solo inventors have been found to develop proportionally more highly-cited patents than those working in association of others (Dahlin, Taylor & Fichman 2004).

At the organizational level, scholars including Henderson (1994), and Iansiti and Clark (1994), have noted the need for firms not only to specialize in disparate domains, but to better integrate their knowledge with a system-wide perspective. If it is difficult for one individual to learn specialist knowledge in multiple domains, it would certainly take much larger effort or longer time for firms to develop such integrative capability. This slowness is borne out by Henderson's (1994) case studies of cardiovascular drug firms.

The preceding discussion leads us to believe that increased trans-specialist experience in innovation projects should have a 'J-shape' effect on the performance of the next project when the trans-specialist knowledge is needed. Note that before the correct causal mechanism for successful commercial performance is understood, there may be situations where an individual may learn the 'wrong' skills for a new app, or inappropriately apply knowledge learned from a previous app in an attempt to improve the commercial performance of a new app (Haleblian & Finkelstein 1999; Finkelstein & Haleblian, 2002; Zollo, 2009). This leads to the following hypothesis:

Hypothesis 2: The commercial performance of successive innovations developed by an individual acting as a trans-specialist has a 'J' shape.

The opposing nature between Hypotheses 1 and 2 can help explain the otherwise no-significant impact when solo and group experiences are aggregated together. For instance, in studying the value of comic books as collectors' items, Taylor and Greve (2006) did not find a significant impact from tenure (naturally correlated with the number of projects involved) for the comic book creators.

Cross-over Learning between Specialist and Trans-Specialists

Innovations are inherently different from one another, and hence both increasingly specialized and integrative knowledge can impact performance. The preceding discussion underlines how the dynamics of the two mechanisms can be very different from each other.

Our discussion has so far carefully crafted the differences in prior experience between the role of a specialist and the role of a trans-specialist. The former represents someone with deep knowledge within one domain, while the latter represents another with at least elementary knowledge in multiple domains. According to the economics of knowledge acquisition described earlier (Postrel, 2002), it would be more costly for the former to acquire knowledge in a hitherto 'foreign' domain than for the latter to start acquiring more specialist domain building on the already-acquired elementary knowledge. If there is any cross-over performance effect from prior innovation experience in these two types of knowledge, switching from being a trans-specialist to a specialist would therefore be much more likely, efficient and effective than the other way round. This presents an asymmetrical dynamic that has so far eluded behavioral scientists.

In other words, prior experience in innovative projects as a trans-specialist in one way or another increase the degree of specialization in certain knowledge domains, and this increase in

specialist knowledge can help in the transition of a focal individual from being a trans-specialist in past innovations to being a specialist (collaborating with other specialist) in future innovations. By definition, patented innovations developed by the inventor of Segway all represented state-of-the-art frontiers in some knowledge domains – so the inventor inadvertently developed deep specialist skills from being a trans-specialist at some point. In contrast, prior experience in innovative projects as a specialist likely does not add to one’s knowledge in a new domain sufficiently substantial for the specialist to acquire further knowledge in that new domain. This leads us to develop the following hypothesis:

Hypothesis 3: The prior experience of an individual acting as a trans-specialist in innovation projects has a positive impact on the commercial performance of future innovation projects in which this individual acts as a specialist.

Empirical Context: Facebook Apps

Few studies document how opportunities are recognized or creatively developed, largely because of a lack of data. The emergence of internet-related platforms and the increasing interest among for-profit companies behind these platforms to encourage the public to develop useful software applications have enabled researchers to track how individual creative innovation is received in the marketplace.

In May 2007, Facebook began to invite individual developers to create applications, or ‘apps’, for Facebook users. Facebook was the first social network site to open its app development platform to the public, and its apps generated much excitement among internet entrepreneurs and venture capitalists (Richmond, 2007). We draw on user-written Facebook applications – or simply ‘apps’ – for our empirical data.

While Facebook apps were downloaded free of charge to users, paid advertisements were just a click away from where apps were downloaded – and the revenue generated from these provided an important motivation for entrepreneurial app developers (Geron, 2007). The allure of reaching a large audience and reaping sizable profits for relatively very little effort encouraged many developers to dabble in app development (Waters, 2010). At least one California-based venture capital firm reportedly raised US \$300 million to develop Facebook apps (Richmond 2007), and many other ventures have followed after our data was sampled (Hagel and Brown 2008). In particular, the user download performance – often in thousands – would be far more difficult for the developers to predict or influence in advance than, say, how efficient their developmental effort would be. Since the downloading and installation of each app has been voluntary for Facebook users, the number of installations, or downloads, of a particular app represents an independent, exogenous measure of its utility or value. Much more so in Facebook apps than in motion pictures, users chose which apps to download based on their functionality or aesthetics, but much less on who develop them.

Moreover, the commercial performance of Facebook apps relies on the perception and adoption by other users. Technical knowledge on how to create Facebook apps therefore constitute only one part of the commercial success of a particular app – knowing how an app might be used, and how potential users may react to it constitute another important component of commercial success. In fact, some technically simple apps have proved to be immensely popular (Richmond 2007). In this sense, the technical knowhow on how to create Facebook apps (itself can be readily divided into different sub-specialties) and the knowhow on what kind of apps may appeal to a wide audience represent at least two knowledge domains or subsystems in developing Facebook apps. Aran and Walker (2011) showed the manner in which Facebook app developers

disseminate information about an app to the friends of adopters of this app has a significant impact on the total user download performance.

Verified by email responses from app developers, the difficulty to predict what apps would be popular required users to think hard about what kind of apps to develop and what these apps are intended to achieve if they were intent on maximizing the commercial potential. Because downloading an app cost time and effort, users were generally not interested in downloading apps that did exactly the same as a prior app. Hence developing replicas of popular apps did not represent a profitable path for developers (consistent with prior studies on software adoption, e.g., Venkatesh & Davis 2000, and Venkatesh et al., 2003). Instead, developers sought to create something with a new functionality or aesthetics to appeal to users

Data: User-contributed Facebook Applications

We collected our data on Facebook apps via several ‘web-crawls’ in 2008 and 2009, essentially containing the information all apps released before that time (since May 2007). After the end of 2008 Facebook stopped reporting the download statistics for each app, and hence only apps released prior to the end of 2008 were used in the analysis. A significant portion of these apps were developed by organized, corporate-like entities, and these showed no information on the actual individuals involved. We excluded these because we could not ascertain the development history of the developers, or simply how many of them were involved. This left us with 32,204 apps developed by 21,746 developers, among whom 18,017 individuals had been involved in only one app, solo or collaboratively, over the observation window.

Since we were interested in the impact of how experience in innovation in the form of Facebook app development affects future performance, we confined our sample to those

developers with at least three apps during our observation period and submitted this cut-off to sensitivity checks. Anecdotal evidence suggested that many users were keen to develop one or two apps as a hobby, but these ‘amateurs’ likely did not dedicate their effort to maximize adoption by other users. On the other extreme, a dozen of so extremely productive individuals developed so many apps within short spans of time that little effect of learning from past experience might be possible. For instance, the most productive developer had 314 apps released in a period of slightly over one year – leaving very little time between successive apps for any learning to take place. To account for how developers ‘learn’ from past innovation experience, we excluded from further analysis 1% of our apps developed by the most prolific individuals (11 in total). The most prolific developer in our eventual sample released fewer than 50 apps during the period. We manually checked the stated functionality of these apps, and noted how they were different from one another (i.e., one is not simply a new version of an old app), ensuring that we indeed worked with creative innovations.

As in patent citations, the performance of these apps was skewed (Griliches, 1981). According to our download statistics on the last day of 2008, 10% of apps accounted for 97% of all download activity, and 5% of apps accounted for 92% of downloads (with an average of 270,626 downloads in this category compared to an overall mean of 16,370).

In the final sample for analysis, there were 7088 apps by 1523 developers. Among these, 803 unique developers (responsible for a total of 4,134 apps) never collaborated with others – we called these ‘habitually solo developers’. In contrast, 374 developers (responsible for 1,218 unique apps) never worked alone in our sample – we called these ‘habitual collaborators’. The rest were 346 developers (responsible for 2,194 unique apps) who sometimes worked alone and sometimes collaborated – we called these ‘mixed’ developers.

By construction, the habitually solo developers took on the role of trans-specialists in developing innovations – we therefore used their learning dynamic to test Hypothesis 2. The habitual collaborators were much more likely to take on the role of specialists than others, and we used their learning dynamic to test Hypothesis 1. The mixed developers were likely trans-specialists in nature but took on the role of specialists in collaborative innovations. We used the solo innovations developed by a mixed developer to supplement the experience of habitually solo developers, and collaborative innovations by a mixed developer to supplement the experience of habitual collaborators. We used the experience of the mixed developers – because of their dual role as trans-specialist in solo innovations and likely specialists in collaborative ones – to test Hypothesis 3.

In general, it is possible that in certain groups, at least one individual may possess all the combined knowledge of the specialists. For instance, authors of certain academic journal articles, especially when they are of the supervisor-supervisee relations, may indeed be able to overcome the coordination challenge in groups of mutually ignorant specialists. However, the preponderance of evidence on how individuals function in groups points to the persistence of communication and coordination challenges that are consistent with collaborative experience among mutually ignorant specialists (Lewis, Lange & Gillis 2005; Liang, Moreland & Argote 1995). We therefore implicitly associated individuals in collaborative innovation projects with the role of specialists.

Econometric Estimation

We constructed the entire app development history for each individual developer. The data structure resembled a panel format – with each developer being a cross-section and the

sequentially released apps a form of time series. This allowed us to exploit analytical tools developed for panel data for econometric estimations. We used STATA xtreg/xtlogit functions to conduct panel data analyses on the data, clustering robust standard errors around each individual developer to account for non-independence in apps developed by the same individuals. In particular, we wanted to control for any early- or late-entry effect for all developers – and therefore included in our controls the number of days elapsed between our first recorded app and the date of release of the first app for a particular developer. Only random-effects model could allow us to include this time-invariant control variable in our model. If this variable was excluded, the Hausman test confirmed the appropriateness of the random effects model – hence this model was adopted.

It was easy to attribute the performance of solo apps to their developers. There was considerably more ambiguity in attributing the performance of collaborative apps to their respective members. In the setting of Facebook apps, reward appropriation and commercial espionage had to be enforced by the developers themselves. For individuals aiming at commercial gains, collaboration worked only to the extent that future gains would not be compromised through the exchange of information required in a group project. As such, it was unlikely that every developer involved would display his or her entire gamut of specialist knowledge in a group project. Rather, one individual likely acted as a ‘leader’ or project integrator in developing a collaborative Facebook app. We therefore tried two methods of performance (and experience) attribution: i) we used the first(-listed) individual of a collaborative app to be the focal developer to whom we attributed performance (and entered the relevant information on prior Facebook app experience), and ii) we randomly selected one of the constituent developers listed under a collaborative app to whom we attributed performance.

Both methods yielded qualitatively similar results, but depending on whether the first individual happened to be a mixed versus habitual collaborators, the number of observations of one group would be inadvertently reduced.

Measures

To measure the commercial performance of a Facebook app as a creative innovation, we used two statistics. The first, called **LnDownloads**, was the logarithm of the total number of downloads for each app – the number of downloads being a close proxy to the advertising revenue a developer might earn. This variable allowed us to gauge whether a developer ‘improved’ from one app to another in terms of achieving more downloads. The large variability in the download performance meant that the logarithmic data more closely followed a normal distribution than the level data. The Facebook reported the total downloads data until the end of 2008 – that determined the end-date of our data collection period, and the date when we took the measure of total downloads.

Following prior studies on breakthrough innovation performance (e.g., Audia & Goncalo, 2007; Singh & Fleming, 2010), we created a second dependent variable: a dummy variable called **KillerApps**, taking on the value of 1 if the focal app was in the top 10 percentile (or more formally, the 90th percentile) of total downloads relative to apps released 15 days before and 15 days after, and 0 otherwise. Because of temporal factors such as seasonality or school holidays, apps developed at roughly the same time were more likely to be the reference group for the focal developers. We tested several different ranges of comparable apps (for instance, 30 days before and after the focal app and the top 5% as the threshold) and our results were not very sensitive to this window selection. Ideally we expected the independent variables to exert similar impact on

both performance measures, but understandably, KillerApps would be more difficult to predict or control from the perspective of the developers.

There are four independent variables in our analysis, designed in the same manner as in prior behavioral studies on how past experience impact future performance (e.g., Taylor & Greve, 2006). **PriorSoloApps** was the total number of apps developed by a developer alone prior to the focal app, with its square term denoted by **PriorSoloAppsSq**. The prior number of solo apps developed by an individual indicated the amount of experience that an individual would be forced to manage multiple knowledge domains or subsystems involved in developing a Facebook app – at least straddling the technical side of development as well as the marketing-oriented intelligence. It would therefore be a reasonable proxy for the amount of trans-specialist experience of a developer.

PriorGroupApps was the total number of apps developed in collaborative groups to which a focal developer belongs, with its square term denoted by **PriorGroupAppsSq**. Individuals working in a collaborative manner typically focused on a subset of the overall problem, and their specialist knowledge developed would be likely sought out by others – affording further opportunities for specialization in the same domain. These two therefore represented the amount of specialist knowledge a developer likely acquired.

A series of control variables were added to highlight the difference between our hypothesized effects and related items from extant theories. First, we included **PriorTeamTies**, which measured the average number of dyadic ties among developers of a collaborative project, to control for the impact of team familiarity on performance. Second, **FirstLag**, indicating the number of days elapsed between the first app in our sample and the first app of the focal developer, was included to control for any early- or late-entry effect for the developers. Third,

Ndayrelease is the number of apps released in a day by a developer. A handful of apps might be developed and released at roughly the same time and it would be hard to tell how these apps influenced each other. Fourth, we included a series of date-related ‘fixed-effect’ dummy variables for apps released within the same 30-day period to account for seasonal patterns of adoption patterns. Eagle *et al.* (2009) documented strong diurnal and weekly patterns of interpersonal communication – directly impacting how Facebook apps were downloaded or spread. Lastly, we created a series of dummy variables to control for the categories of an app. Each developer could designate up to two app categories for each app, although designating an app category was not mandatory – and many apps did not have a designated category.

We tried to control for the duration for which an app had been released (**Duration**, indicating the number of days a focal app had been available for downloads). However, including this variable increased the risk of multi-collinearity without significantly changing the coefficients of our independent variables. We believe that the effect of release duration on total downloads was eclipsed by the seasonal effect of when an app was released – this latter effect was captured in our use of dummy variables representing a 30-day period within which apps were released. We therefore did not include the Duration variable in our reported results. To control for the ‘supply’ of and ‘demand’ for apps at any particular time: **AppTotalSupply** measured the cumulative sum of all Facebook apps (in thousands) released prior to a focal app, while **AppTotalDemand** measured the cumulative total downloads (in millions) undertaken prior to a focal app. However, these additional control variables were multi-collinear with the time fixed-effects dummies and did not significantly improve the fitness model fitness. We therefore did not incorporate these variables in our models.

In addition, to facilitate a comparison between our findings with those from standard behavioral learning, we included an estimation using an aggregate measure of past experience for mixed developers: **PriorApps** – denoted as the total number of apps developed by a developer prior to a focal app, and its squared term **PriorAppsSq**.

Estimation Results

The means, standard deviations and inter-correlations for the main variables included in the econometric analysis for the entire sample are shown in Table 1. For a comparison, Table 2 shows the descriptive statistics for different categories of app developers.

Table 3 shows the random-effects analyses on the commercial performance of Facebook apps developed by habitually solo developers. For the regression in this and following tables, we also ran the OLS regression for each model and to ensure that multi-collinearity was not a serious issue. Model 3a shows the controls-only model. Model 3b shows the model with only the linear term **PriorSoloApps**, with an estimated coefficient that was not statistically significant. When the quadratic term was included in Model 3c, the estimated coefficients for both **PriorSoloApps** and **PriorSoloAppsSq** became statistically significant ($p < 0.05$ and $p < 0.01$ respectively). The negative and positive signs of these two variables confirm Hypothesis 2 – while rejecting a simple linear relation between past solo experience and download performance. It took about 30 apps, however, for the total download performance to reverse the initial course of decline and start improving. A similar pattern was observed in Model 3d when **KillerApps** was used as the dependent variable, and the estimated coefficients for both **PriorSoloApps** and **PriorSoloAppsSq** were marginally significant at $p < 0.1$.

Table 4 shows the random-effects analyses on the commercial performance of Facebook apps developed by habitually collaborative developers. Models 4a through 4d used **LnDownloads** as the dependent variable, with Models 4a through 4c using random developer performance attribution. Model 4a shows the controls-only model having randomly attributed the performance of a collaborative app to any one of the constituent member. Model 4b added the linear term **PriorGroupApps** to the analysis, and its estimated coefficient was not statistically significant. Model 4c added the quadratic term **PriorGroupAppsSq**. The estimated coefficients for the linear and quadratic terms in Model 4c were statistically significant ($p < 0.01$ and $p < 0.05$ respectively), and with their appropriate signs, confirmed Hypothesis 1. Model 4d showed the same analysis while attributing the performance of the collaborative app to the first-listed developer – quantitatively similar to that of Model 4c. Models 4e and 4f used **KillerApps** as the dependent variable. Model 4e used a random performance attribution for the collaborative apps, and the estimated coefficients for the linear term **PriorGroupApps** and the quadratic term **PriorGroupAppsSq** were respectively negative and positive in sign, but only marginally significant ($p < 0.1$) – providing partial support for Hypothesis 1. In Model 4f, where app performance was attributed to the first-listed developer neither coefficient was statistically significant (both took on a negative sign).

Table 5 shows the random-effects regressions on the performance of only the collaborative apps developed by mixed developers. Models 5a through 5g used **LnDownloads** as the dependent variable, with Models 5a through 5e using random developer performance attribution. Model 5a shows the controls-only model having randomly attributed the performance of a collaborative app to any one of the constituent member. Model 5b added the **PriorApps** (positive at $p < 0.05$) and Model 5c added the **PriorAppsSq** (negative at $p < 0.01$)

variables. Aggregating all prior group and solo app experiences therefore yielded coefficients that mirrored the traditional learning curve with performance increasing at a decreasing rate. The similar inverted U-shaped impact of **PriorApps** seems to suggest that it may be driven by the prior group experience.

Model 5d investigated the separate impact of prior group and prior solo experience with the linear terms – the linear terms generally yielded a poor fit. Model 5e added the quadratic terms of the prior experience. The curvilinear impact of **PriorGroupApps** (both linear and quadratic terms at $p < 0.05$, with the appropriate signs) again confirmed Hypothesis 1. The positive and statistically significant ($p < 0.05$) impact of **PriorSoloApps** confirmed Hypothesis 3. Models 5f and 5g relied on performance attribution to the first-listed developers, but neither showed significant support to our hypotheses – possibly because of the dramatically reduced sample size (first-listed developers had to be in the mixed category). Models 5h and 5i used **KillerApps** as the dependent variables but while the estimated coefficients were consistent with the curvilinear impact of **PriorGroupApps**, they were not statistically significant. Since **PriorApps** was simply the sum of **PriorGroupApps** and **PriorSoloApps**, all three variables could not be included in the same analysis.

Separately, we also tried using a fixed-effect system where dummy variables represent individual developers involved in a particular group project, but this reaction generated very poor fit overall. In an additional analysis, we did check to see that the performance of solo apps developed by mixed developers was not significantly influenced by the amount of prior collaborative experience (i.e., **PriorGroupApps** and **PriorGroupAppsSq**) of the focal developer. Together, this is consistent with our argument that there is an asymmetrical dynamic

in different experience and it is more likely for trans-specialists to help the specialists for better innovation performance rather than the other way around.

Discussion and Conclusion

This study was motivated by a fundamental question in whether individuals could learn to develop better innovations – an idea that has long intrigued psychologists and more recently organization scientists. Researchers who have hoped to find a simple equivalent to the traditional learning curve derived from manufacturing contexts have been clearly frustrated by the lack of overall significant relations: that past experience has not until now been found to have a significant impact on the commercial performance of future innovations (e.g., Taylor & Greve, 2006). The continued lack of significance in this finding questions some long-held assumptions in knowledge acquisition, including the conventional wisdom in its path-dependency (Penrose, 1959; Cohen & Levinthal, 1990; Nelson & Winter, 1982; Zander & Kogut, 1995).

We began as our first principles the often cited findings that creativity relies on a diversity of knowledge expertise rather than mere information scanning or exposure (Taylor & Greve, 2006; Tripsas, 1997). Our fresh theoretical perspective based on specialist versus trans-specialist orientations in individuals disentangles the solo versus group innovation experience an individual may accrue, and thereby identifies two different ‘learning curves’ for the different experiences.

We found in the Facebook applications development context that group work experience may intensify the individuals’ tendency for specialization in sub-problems, and the performance improvement was at first very positive and significant, but incremental improvement took place at a decreasing rate – eventually leading to performance declines. We also found that once

individual inventors were able to internalize the specialization in different knowledge domains and develop an integrative capability across knowledge domains, the payoff would be positive in the long run, as suggested by the positive impact of prior solo works experience – although the manifestation of such ‘learning’ would take considerable time.

Our finding suggests that the long-held assumption of path-dependency may still hold in the arena of innovation development, only that it takes on a somewhat unfamiliar shape, and that the specialist or trans-specialist nature of the experience makes a difference. At the same time, the curvilinear shapes confirmed by our empirical analysis throw into question whether more experience is definitely better in terms of innovation experience. In the context of innovations – where creativity is required, it is no longer so simple as to suggest that more experience would definitely lead to better commercial performance. The (specialist versus trans-specialist) nature of that experience, together with the counts of experience, matters. In this manner, our study definitely extends the traditional learning curve literature into the innovation domain (e.g., Argote & Epple, 1990). In fact, the traditional learning curve derived from more repetitive, manufacturing-oriented settings forms an important foundation of our theoretical perspective: that increasingly specialized knowledge also improves the possibility of an innovation – what it can deliver, how, and at what quality. From the perspective of an individual programmer, more collaborative development within the same Facebook context further the specialist knowledge – which is useful whether it is with the same collaboration partners. It is only in the trans-specialist (or system-wide) perspective where the learning curve effect is impeded (e.g., Williams & O’Reilly, 1998). Scholars are therefore urged to pay close attention to how the traditional form of learning accrues in complicated innovation systems.

Another important question in managing creativity in organization is how best to utilize the experience of organizational members for future team design or creative tasks assignment, and whether this would differ for team experience and individual experience (Taylor & Greve, 2006: 723). Our empirical results suggest that different kinds of experience will have different impact on the capability development in individuals and therefore influence how the organization is ought to combine these capabilities from different individuals. We concur with the suggestion (Taylor & Greve, 2006) that when managers staff cross-functional, cross-knowledge team for creative tasks, it is essential that the included members have deep understandings of their respective knowledge domains. To let individuals participate in various groups may help them to exploit and deepen their specialization in a particular domain. But if the organization plans to cultivate a “super” talent to lead in seeking innovation, our results suggest that a combination of both solo work and collaborative work are required. This is exactly in the spirit of March’s (1991) suggestion on personnel turnover to maintain the creative capabilities.

The current research also has several limitations that could be overcome by the future research. First of all, although we track the collaboration history of inventors, we could not track the exact extent of knowledge diversity within an individual. How loosely our measure is linked to these underlying knowledge evolution process calls for caution in interpreting our results beyond our context. Although most studies might be limited by the data constraints in examining at this deep level, some smart experimental design (e.g, Girotra et al., 2010 using hybrid structure) might help us to alleviate the concern. Second, our econometric analysis focus on the mean performance, we do not explicitly concern about variance in our theoretical development, although this itself can be a fruitful avenue for further research (e.g., Singh & Fleming, 2010; Taylor & Greve, 2006). Third, we could not control for the application content

beyond application categories due to the data limitation. It was possible that the nature of the application or task context creates a trade-off between different aspects of learning or specialization in product development (Taylor & Greve, 2006). To study the specific knowledge domains, such as the categories in which apps were assigned, further research in the development of these apps would be useful.

Finally, it is also important to examine the group composition which is crucial for the experience acquisition in collaborative creation. The empirical implication of our theoretical arguments on system-wide perspective may be contingent on what experience can be learned from the group as a whole and from the specific dyads interaction between two group members. Our empirical support will be strengthened if the group outcome is influenced by the former than the later.

REFERENCES

- Ahuja, G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Administrative Science Quarterly*, 45: 425–455.
- Amabile, T. M. 1996. *Creativity in Context*. Boulder: Westview.
- Aran S, Walker D. 2011. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science* 57(9), 1623-1639.
- Åstebro T, & Thompson P. 2011. Entrepreneurs, Jacks of all trades or Hobos? *Research Policy*, 40(5): 637-649.
- Audia P. G. , & Goncalo, J. A. 2007. Past success and creativity over time: A study of inventors in the hard disk drive industry. *Management Science*, 53(1), 1-15.

- Argote, L., & Epple, D. 1990. Learning curves in manufacturing. *Science*, 247:920-924.
- Bardolet, D., Fox, C. R., & Lovallo, D. 2011. Corporate capital allocation: A behavioral perspective. *Strategic Management Journal*, 32(13): 1465-1483.
- Barkema, H. G. & Schijven, M. 2008. How do firms learn to make acquisitions? A review of past research and future agenda for the future. *Journal of Management*, 34(3), 594-634.
- Benner, M. J., & Tushman, M. L., 2002. Process management and technological innovation: A longitudinal study of the photography and paint industries. *Administrative Science Quarterly*, 47: 676-706.
- Cohen, W. M., & Levinthal, D. A. 1990. Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly*, 35:128-152.
- Crossan, M. M., & Apaydin, M. 2010. A multi-dimensional framework of organizational innovation: A systematic review of the literature. *Journal of Management Studies*, 47(6): 1154–91.
- Csikszentmihalyi, M. 1988. Society, culture and person: A system view of creativity, in Sternberg, R. J. (Ed), *The Nature of Creativity*. New York: Cambridge University Press.
- Cummings, J. N. 2004. Work groups, structural diversity, and knowledge sharing in a global organization. *Management Science*, 50(3): 352–64.
- Cyert, R. M., & March, J. G. 1963. *A Behavioral Theory of the Firm*. Prentice Hall, Englewood Cliffs, NJ.
- Dahlin, K., Taylor, M., & Fichman, M. 2004. Today's Edisons or weekend hobbyists: Technical merit and success of inventions by independent inventors. *Research Policy*, 33(8): 1167-83.

- Darr, E.P., Argote, L., & Epple, D. 1995. The acquisition, transfer and depreciation of knowledge in service organizations: Productivity in franchises. *Management Science*, 41(11):1750-1762.
- Denrell, J. 2003. Vicarious learning, undersampling of failure, and the myths of management. *Organization Science*, 14(3): 227–243.
- Eagle N, Pentland A, Lazer D. 2009. Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences of the United States of America*. 106 (36), 15274-15278.
- Ferriani, S., Cattani, G & Baden-Fuller, C. 2009. The Relational Antecedents of Project-Entrepreneurship: Network Centrality, Team Composition and Project Performance. *Research Policy*, 38 (10): 1545–1558.
- Fleming L. 2000. Recombinant uncertainty in technological search. *Management Science*, 47(1): 117-132.
- Fleming, L., Santiago, M., & Chen, D. 2007. Collaborative brokerage, generative creativity, and creative success. *Administrative Science Quarterly*, 52: 443–75.
- Geron, T. 2007. Financing: Stanford’s Facebook apps class proves a quick hit. *Wall Street Journal (Eastern online edition)*, 18 December.
- Girotra, K., Terwiesch, C., & Ulrich, K. T. 2010. Idea generation and the quality of the best idea. *Management Science*, 56(4): 591-605.
- Griliches, Z. 1981. Market value, R&D, and patents. *Economics Letters*, 7:183-187.
- Hagel, J., & Brown, J. S. 2008. Life on the edge: Learn from Facebook.
http://www.businessweek.com/innovate/content/apr2008/id2008042_809134.htm.
(Accessed December 2009)

- Haleblian, J., & Finkelstein, S. 1999. The influence of organizational acquisition experience on acquisition performance: A behavioral learning perspective. *Administrative Science Quarterly*, 44(1): 29–56.
- Harris M. 2010. What comes first? *Engineering and Technology*, 4(21): 28-29.
- Haunschild, P. R., & Sullivan, B. N. 2002. Learning from complexity: Effects of prior accidents and incidents on airlines' learning. *Administrative Science Quarterly*, 47:609–643.
- He, Z. L. & Wong, P. K., 2004. Exploration vs. exploitation: An empirical test of the ambidexterity hypothesis. *Organization Science*, 15: 481-494.
- Henderson R. 1994. The evolution of integrative capability: Innovation in cardiovascular drug discovery. *Industrial and Corporate Change*, 3: 607-630.
- House, R., Rousseau, D. M., & Thomas-Hunt, M., 1995. The meso paradigm: A framework for the integration of micro and macro organizational behavior. *Research in Organizational Behavior*, 17: 71-114.
- Hsieh, C., Nickerson, J. A., & Zenger, T. R. 2007. Opportunity discovery, problem solving and a theory of the entrepreneurial firm. *Journal of Management Studies*, 44(7): 1255-1277.
- Iansiti, M., & Clark, K. B. 1994. Integration and dynamic capability: Evidence from product development in automobiles and mainframe computers. *Industrial & Corporate Change*, 3(3) 557-605.
- Kim, J-Y., Haleblian, J., & Finkelstein, S. 2011. When firms are desperate to grow via acquisition: The effect of growth patterns and acquisition experience on acquisition premiums. *Administrative Science Quarterly*, 56(1): 26-60.

- Kim J.-Y., & Miner, A. S. 2007. Vicarious learning from the failures and near-failures of others: evidence from the U.S. commercial banking industry. *Academy of Management Journal*, 50(3): 687-714.
- Kim, J.Y., Kim, J. Y., & Miner, A. S. 2009. Organizational learning from extreme performance experience: The impact of success and recovery experience. *Organization Science*, 20(6): 958-978.
- Kirton, M., 1994. *Adaptors and Innovators: Styles of Creativity and Problem Solving*, Routledge, New York
- Lampel, J., Lant, T., & Shamsie, J. 2000. Balancing act: Learning from organizing practices in cultural industries. *Organization Science*, 11(3): 263–9.
- Levinthal, D., A., & March, J. G. 1993. The myopia of learning. *Strategic Management Journal*, 14: 95-112.
- Lewis, K., Lange, D., & Gillis, L. 2005. Transactive memory systems, learning, and learning transfer. *Organization Science*, 16:581–598.
- Liang, D., Moreland, R., & Argote, L. 1995. Group versus individual training and group performance; The mediating effects of transactive memory. *Personality Soc. Psych. Bull*, 21: 384–393.
- Lingo, E. L., & O'Mahony, S. 2010. Nexus work: Brokerage on creative projects. *Administrative Science Quarterly*, 55(1): 47-81.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science*, 2: 71–87.
- March, J. G., Sproull, L. S., & Tamuz, M. 1991. Learning from samples of one or fewer. *Organization Science*, 2(1): 1–14.

- Moore, P. 2006. Microprocessor scaling: What limits will hold? *IEEE Annals of the History of Computing*, 28(3): 89.
- Nelson, R. R., & Winter, S. G. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press.
- Obstfeld, D. 2005. Social networks, the *Tertius iungens* orientation, and involvement in innovation. *Administrative Science Quarterly*, 50 (1): 100-130.
- Parker, S. C. 2006. Learning about the unknown: How fast do entrepreneurs adjust their beliefs? *Journal of Management Studies*, 21 (1): 1-26.
- Penrose, E. 1959. *The Theory of the Growth of the Firm*. Oxford University Press, England, UK.
- Perry-Smith, J. E. 2006. Social yet creative: The role of social relationships in facilitating individual creativity. *Academy of Management Journal*, 49(1): 85-101.
- Perry-Smith, J. E., & Shalley, C. E. 2003. The social side of creativity: A static and dynamic social network perspective. *Academic of Management Review*, 28(1): 89-106.
- Perry-Smith, J. E., & Coff, R. W. 2011. In the mood for entrepreneurial creativity? How optimal group affect differs for generating and selecting creative ideas for new ventures. *Strategic Entrepreneurship Journal*, in press.
- Postrel, S. 2002. Islands of shared knowledge: Specialization and mutual understanding in problem-solving teams. *Organization Science*, 13(3): 303-320.
- Ratliff, E. 2006. Me translate pretty one day. *Wired*, 14(12): 210-213.
- Reagans, R., Argote, L., & Brooks, D. 2005. Individual experience and experience working together: Predicting learning rates from knowing who knows what and knowing how to work together. *Management Science*, 51(6): 869-881.

- Reagans, R., & McEvily, B. 2003. Network structure and knowledge transfer: The effects of cohesion and range. *Administrative Science Quarterly*, 48(2): 240-267.
- Richmond, R. 2007. Enterprise: Why so many want to create Facebook applications: site's growing ranks seen as potential source of revenue, customers. *Wall Street Journal* (Eastern edition), 4 September, B4.
- Schilling, M.A., Vidal, P., Ployhart, R. E., & Marangoni, A. 2003. Learning by doing something else: Variation, relatedness, and the learning curve. *Management Science*, 49(1): 39-56.
- Schwartz, E. I. 2004. Sparking the fire of invention. *Technology Review*, 107(4): 32-40.
- Simonton, D. 1999. *Origins of genius: Darwinian perspectives on creativity*. New York: Oxford.
- Singh, J., & Fleming, L. 2010. Lone inventors as sources of breakthroughs: Myth or reality? *Management Science*, 56(1): 41-56.
- Sorensen, J., & Stuart, T. E., 2000. Aging, obsolescence, and organizational innovation, *Administrative Science Quarterly*, 45: 81-112.
- Sosa, M. E. 2011. Where do creative interactions come from? The role of tie content and social networks. *Organization Science*, 22(1): 1-21.
- Taylor, A., & Greve, H. R. 2006. Superman or the fantastic four? Knowledge combination and experience in innovative teams. *Academy of Management Journal*, 49(4): 723-40.
- Tripsas, M. 1997. Surviving radical technological change through dynamic capability: Evidence from the typesetter industry. *Industrial and Corporate Change*, 6: 341-377.
- Thornton, P. 2002. The rise of the corporation in a craft industry. *Academy of Management Journal*, 45: 81-101.

- Venkatesh, V., & Davis, F. D. 2000. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science*, **46(2)**: 186-204.
- Venkatesh, V., Morris, M. G., Davis, G. B. & Davis, F. D. 2003. User acceptance of information technology: Toward a unified view. *MIS Quarterly*, **27(3)**: 425-478.
- Wang, L., & Zajac, E. J. 2007. Alliance or acquisition? A dyadic perspective on interfirm resource combinations. *Strategic Management Journal*, **10**: 307-322.
- Waters, R., 2010. A code to crack. *Financial Times*, 17-18 April, Asia Edition: 8.
- Williams, K. Y., & O'Reilly, C. A. 1998. Demography and diversity in organizations: A review of 40 years of research. In B. M. Staw & R. I. Sutton (Eds.), *Research in organizational behavior*, vol. 20: 77–140. Greenwich, CT: JAI Press.
- Yayavaram, S., & Ahuja, G. 2008. Decomposability in knowledge structures and its impact on the usefulness of inventions and knowledge-based malleability. *Administrative Science Quarterly*, **53(2)**: 333-362.
- Zaheer, A., & Soda, G. 2009. Network evolution: The origins of structural holes. *Administrative Science Quarterly*, **54(1)**: 1-31.
- Zander, U., & Kogut, B. 1995. Knowledge and the speed of the transfer and imitation of organizational capabilities: An empirical test. *Organization Science*, **6**: 76-92.
- Zollo, M. 2009. Superstitious Learning with Rare Strategic Decisions: Theory and Evidence from Corporate Acquisitions. *Organization Science*, **20 (5)**: 894-908.

Table 1: Descriptive statistics for main variables in the regression analysis for all observations (i.e., no developers are randomly selected out)

	Mean	S.D.	Min	Max	1	2	3	4	5	6	7	8
1 LnDownloads	4.32	3.45	0	15.93	1.00							
2 KillerApps	0.12	0.33	0	1	0.64	1.00						
3 PriorSoloApps	2.57	4.99	0	44	-0.20	-0.10	1.00					
4 PriorGroupApps	2.15	5.30	0	44	-0.06	0.05	-0.13	1.00				
5 PriorApps	4.72	6.79	0	44	-0.20	-0.03	0.63	0.68	1.00			
6 FirstLag	272.00	155.00	24	563.00	-0.59	-0.17	0.10	-0.07	0.02	1.00		
7 Applapse	21.46	47.72	0	531.00	0.02	0.09	-0.06	-0.04	-0.07	-0.17	1.00	
8 Ndayrelease	1.81	2.56	1	24	0.10	0.15	-0.01	0.00	0.00	0.03	0.25	1.00
9 PriorTeamTies	1.75	4.71	0	40	-0.01	0.10	-0.16	0.87	0.56	-0.04	0.04	0.23

7088 apps for 1523 developers

Table 2: Within-developer descriptive statistics

	Habitually Solo Developers		Habitually Collaborative Developers		Mixed Developers	
	Mean	std. dev	Mean	std. dev	Mean	std. dev
PriorSoloApps ⁽¹⁾	3.96	4.84			2.91	3.9
PriorGroupApps			4.91	6.2	2.7	4.67
LnDownloads ⁽²⁾	3.65	2.68	5.05	2.93	4.99	2.51
Proportion of apps being killer apps ⁽³⁾	6.70%		20.80%		14.80%	
No. of developers	803		374		346	
No. of apps	4134		1218		2194	

(1) First, the maximum value of the variable for a particular developer over the observation window was calculated, then the mean and the standard deviation was calculated across all developers

(2) First, the mean value of the variable for a particular developer over the observation window was calculated, then the mean and the standard deviation was calculated across all developers

(3) The proportion of apps having accomplished the 10% total user downloads among all apps released 15 days before and after

Table 3: Effects of prior solo creative experience in habitually solo developers

VARIABLES	DV: LnDownloads			DV: KillerApps	
	a	b	c	d	e
PriorSoloApps		0.000 [0.006]	-0.036** [0.013]	-0.026 [0.035]	-0.097+ [0.057]
PriorSoloAppsSq			0.001** [0.000]		0.004+ [0.003]
FirstLag	0.002** [0.001]	0.002** [0.001]	0.001* [0.001]	0.004 [0.003]	0.004 [0.003]
Applapse	-0.000 [0.001]	-0.000 [0.001]	-0.000 [0.001]	0.002 [0.004]	0.002 [0.004]
Ndayrelease	-0.026 [0.021]	-0.025 [0.022]	-0.023 [0.022]	-0.201+ [0.104]	-0.212* [0.105]
Constant	6.301*** [0.287]	6.300*** [0.288]	6.297*** [0.288]	-3.061*** [0.778]	-3.042*** [0.781]
F test for fixed time effect	2398.465	2068.519	2064.812	40.349	40.720
F test for fixed category effect	79.676	79.615	78.969	26.929	26.425
Log-likelihood	-7688.597	-7688.597	-7683.383	-782.774	-781.557
sigma_u	1.100	1.100	1.102	1.659	1.677
rho	0.394	0.394	0.396	0.456	0.461

4134 apps by 803 developers

Standard errors in brackets

*** p<0.001, ** p<0.01, * p<0.05, + p<0.1

Table 4: Effect of prior collaborative creative experience in habitually collaborative developers

VARIABLES	DV: LnDownloads				DV: KillerApps	
	"Random" developer			First developer	"Random" developer	First developer
	a	b	c	d	e	f
PriorGroupApps		0.045 [0.035]	0.122* [0.052]	0.091+ [0.047]	0.279+ [0.168]	-0.031 [0.191]
PriorGroupAppsSq			-0.002* [0.001]	-0.002** [0.001]	-0.012+ [0.007]	-0.006 [0.006]
FirstLag	-0.002 [0.002]	-0.001 [0.002]	-0.000 [0.002]	0.001 [0.002]	-0.009+ [0.005]	-0.004 [0.004]
Applapse	0.000 [0.002]	0.001 [0.002]	0.001 [0.002]	0.004+ [0.002]	-0.004 [0.005]	0.002 [0.005]
PriorTeamTies	0.018 [0.018]	-0.018 [0.033]	-0.031 [0.034]	0.010 [0.034]	-0.045 [0.120]	0.237 [0.156]
Ndayrelease	0.252*** [0.044]	0.277*** [0.048]	0.273*** [0.048]	0.267*** [0.046]	0.310* [0.124]	0.372* [0.147]
Constant	6.983*** [0.767]	6.930*** [0.767]	6.896*** [0.765]	5.764*** [0.777]	-2.232 [1.406]	-4.342* [1.696]
F test for fixed time effect	197.512	183.949	177.383	235.387	12.162	21.326
F test for fixed category effect	78.008	79.163	81.379	75.269	24.337	25.373
Observations	721	721	721	876	721	876
Number of developers	146	146	146	237	146	237
Log-likelihood	-1551.491	-1550.649	-1548.652	-1850.439	-239.864	-284.394
sigma_u	1.904	1.904	1.889	1.825	2.294	2.707
rho	0.550	0.551	0.548	0.553	0.615	0.690

Standard errors in brackets

*** p<0.001, ** p<0.01, * p<0.05, + p<0.1

Table 5. Effect of prior creative experience in mixed developers

VARIABLES	DV: LnDownloads						DV: KillerApps				
	"Random" developer					First developer		"Random" developer		First developer	
	a	b	c	d	e	f	g	h	i		
PriorApps		0.039+	0.108***			0.080					
		[0.021]	[0.032]			[0.052]					
PriorAppsSq			-0.003**			-0.003					
			[0.001]			[0.002]					
PriorSoloApps				0.067+	0.127*		-0.019	0.021		0.001	
				[0.039]	[0.062]		[0.105]	[0.200]		[0.384]	
PriorSoloAppsSq					-0.004		-0.001	-0.003		-0.010	
					[0.003]		[0.004]	[0.014]		[0.047]	
PriorGroupApps				0.022	0.078*		0.108+	0.267*		0.292+	
				[0.029]	[0.038]		[0.060]	[0.135]		[0.157]	
PriorGroupAppsSq					-0.003*		-0.003	-0.011+		-0.013*	
					[0.001]		[0.002]	[0.006]		[0.006]	
FirstLag	-0.002+	-0.001	-0.001	-0.001	-0.001	0.001	0.001	-0.002		-0.003	
	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]	[0.002]	[0.002]	[0.004]		[0.005]	
Applapse	-0.003+	-0.003	-0.002	-0.003	-0.002	-0.002	-0.002	0.000		0.004	
	[0.002]	[0.002]	[0.002]	[0.002]	[0.002]	[0.003]	[0.003]	[0.005]		[0.007]	
Ndayrelease	-0.024	-0.000	-0.027	-0.011	-0.039	0.055	0.058	-0.017		-0.106	
	[0.064]	[0.065]	[0.065]	[0.066]	[0.067]	[0.106]	[0.108]	[0.215]		[0.238]	
PriorTeamTies	-0.026	-0.058*	-0.029	-0.042	-0.014	-0.013	-0.035	0.055		-0.045	
	[0.020]	[0.027]	[0.029]	[0.032]	[0.034]	[0.041]	[0.046]	[0.090]		[0.105]	
Constant	7.935***	7.900***	7.908***	7.884***	7.910***	7.074***	7.191***	-2.818*		-2.123	
	[0.488]	[0.487]	[0.484]	[0.487]	[0.486]	[0.693]	[0.698]	[1.137]		[1.425]	
F test for fixed time effect	317.629	308.648	318.434	308.320	316.188	113.215	103.214	12.691		7.332	
F test for fixed category effect	69.250	69.815	71.663	70.172	71.683	38.231	37.576	21.382		11.949	
Observations	758	758	758	758	758	321	321	758		321	
Number of developers	313	313	313	313	313	136	136	313		136	
Log-likelihood	-1583.901	-1582.242	-1578.211	-1581.857	-1577.995	-678.101	-677.523	-215.579		-110.918	
sigma_u	1.547	1.547	1.507	1.552	1.512	1.545	1.549	2.853		1.735	
rho	0.491	0.492	0.479	0.495	0.481	0.467	0.470	0.712		0.478	

Standard errors in brackets

*** p<0.001, ** p<0.01, * p<0.05, + p<0.1

Figure 1. Incremental Cost in Acquiring More Knowledge in a Representative Domain

