Product Updates: Attracting New Consumers Versus Alienating Existing Ones

Jens Foerderer and Armin Heinzl[1]

[1] Business School, University of Mannheim, Germany.
{foerderer,heinzl}@uni-mannheim.de

Abstract. Are product updates—in terms of producers' decision to add new features "over-the-air"—an effective means in stimulating greater product demand and appeal? Our difference-in-differences analyses of a matched sample of 17,247 mobile apps in Google Play over a period of 24 weeks documents mixed consumer reactions. Whereas updates attracted new consumers, existing consumers rated an app 1.1% worse than before the update and compared to a control group of not updated apps. Why did existing consumers react negatively to updates? Our data provides little support that economic reasoning—i.e., direct costs or learning costs imposed by the update—underlies these reactions. Instead, negative reactions appear—at least to a certain degree— as a behavioral phenomenon, as consumer reviews show an increased density of affective vocabulary after an update. We conclude that updates may stimulate new demand but may alienate existing consumers.

Keywords: Software update, Mobile apps, behavioral economics, difference-in-differences, propensity score matching.

1 Introduction

Product updates have become an essential instrument in firms' product management repertoire. Once representing solely a means of providing bug fixes and minor improvements, updates are increasingly being used to provide consumers with new features over-the-air [1–3]. Such "feature updates" allow producers to add new functionality after a product's market release and while it is in use by consumers [4, 5]. Since 2008, for example, Apple has added several hundreds of features to its iPhone, including the "Siri" personal assistant, "facetime" video calls or "iCloud", a feature to synchronize files across devices. The increasing use of software makes feature updates also relevant for many products that traditionally relied little on information technology. For example, cars are increasingly software-based, enabling automakers to push substantial new features to their fleet without requiring consumers to bring their vehicle into the dealer's garage [1, 6]. Tesla, for example, recently made a product update to improve driving performance [7]. Even beyond cars and phones, any other products have become updateable, including televisions, washing machines, or books [3]. In sum, updates enable producers to alter a product in use by consumers, instead of being limited to introducing features only over subsequent new product releases. In this

Multikonferenz Wirtschaftsinformatik 2018,
March 06-09, 2018, Lüneburg, Germany

sense, updates offer producers new possibilities in creating long-lasting demand and appeal for their products [1, 8], eventually representing “a new set of strategic choices related to how value is created and captured” [2].

In this paper, we study the effects of updates on product demand and ratings. Updates are not trivial to study because they affect two distinct audiences: *existing* consumers of a product as well as potentially *new* consumers attracted by an update. Updates may be an effective means to increase demand for a product by attracting new consumers. As updates introduce new features, they may increase the likelihood of a product to enter the consideration set of consumers. This presumption stems foremost from traditional utility theory, which suggest that each additional product attribute that consumers perceive positively increases consumers’ utility [9, 10]. First empirical evidence seems to confirm this suggestion. In the context of browser add-ons, Tiwana [5] finds frequent updates linked to higher downloads. Similarly, consumers seem to choose digital products that offer more features over comparable ones with fewer features [11–14].

Reactions of existing consumers are less clear. Tiwana et al. [5] find a link between updates and higher product ratings. Fleischmann et al. [4] observe in an experimental setting that updates positively disconfirm existing consumers, leading to a greater intention to continue using a product. These contributions notwithstanding, there is anecdotal and empirical evidence that existing consumers react negatively. Anecdotes of product failures suggest for instance a “feature creep”, meaning that the ongoing addition of features may result in bloated and over-complicated products, ultimately making consumers abandon a product [15]. In addition, simple Google searches yield hundreds of tutorials for how to undo recent updates of Instagram, Facebook, or Snapchat, to only name a few. Research so far suggests both rational and behavioral drivers for negative reactions of existing consumers. A rational reason might be that updates confront consumers with costs for switching from the product they know to the updated one [16–19]. Switching costs include transaction costs for an update (e.g., update fees, data plans) and learning costs (e.g., handling new features, changes in the user interface). If switching costs imposed by an update exceed potential utility, consumers may react negatively to the update. Apart from rational considerations, negative reactions may also reflect a behavioral phenomenon driven by psychological ownership [20, 21], reluctance toward novelty [22], endowment [23, 24], or routine-seeking behavior [25]. For example, consumers may be reluctant toward novel things such as new features regardless of potential utility. In sum, however, conclusive empirical evidence on the effect of updates as well as the mechanisms driving these effects remains scarce.

Studying the effects of updates is difficult because it requires detailed product-level data that allows identifying changes made by updates and isolating the reactions of existing consumers and new consumers. An investigation should also account for unobservable and observable differences among products and producers that may confound consumer reactions. Finally, the above considerations underscore that such an investigation requires identifying causality. The decision to update is endogenous to producers, and likely to suffer from reverse causality. Because of this complexity, the accurate measurement of the consequences of updates has proven elusive.

Our study uses data on mobile apps, which allows us to track product-level information on consumer reactions (i.e., ratings, downloads, and reviews) as well as textual information on updates released by app developers (i.e., changelogs). Our dataset comprises a weekly balanced panel of 17,247 distinct apps listed in the Google Play Store U.S., the largest market for apps worldwide, over 24 weeks in 2016. To address endogeneity bias, we use propensity score matching [26, 27]. We match apps that are updated (treatment group) to apps that are not updated (control group) but equivalent given our observational data. In our context, we have access to several thousands of potential control observations that we can use as inputs for the matching procedure, which bolsters our control group design [27, 28].

2 Theoretical Background

Updates are "self-contained modules of software that are provided to the user for free in order to modify or extend a software after it has been rolled out and is already in use" [4]. Updates are no stand-alone products but rather are integrated into the base product [29, 30]. In this paper we are interested in product updates that add new functionalities to a product. We refer to such updates as "feature updates" and use the term "update" for the sake of brevity. An example for feature updates are Apple's regular updates for its iPhone. The 2016 update for instance brought more than 50 new features.

Updates are evaluated by two distinct audiences, which also require differentiated theoretical considerations. One evaluating audience consists of *new* consumers attracted by an update. The other evaluating audience encompasses *existing* consumers of a product. The main argument of our paper is that new consumers and existing consumers react differently to updates. In the following, we discuss this argument in more detail.

Feature updates may attract new consumers by making it more likely that the product enters the consideration sets of consumers. This presumption comes from traditional utility theory, which has modeled consumer preferences using an additive utility function [10]. Utility models assume each additional product attribute that consumers perceive positively to increase consumers' utility. This idea finds itself implemented in many market research techniques, such as the conjoint analysis or discrete choice models [e.g., 9]. Because these models predict consumer outcomes based on expected utilities or part-worths for each product feature, the conclusion is that each positively valued feature adds to the success of a product, compared to not having the feature.

Empirical evidence supports the prediction that consumers choose products that offer more features over comparable products with fewer features [11, 12, 14]. One consistent observation is that adding attributes to a product increases consumers' perceptions of its capability, resulting in improved product evaluations before ownership [12]. Consequently, feature updates are likely to attract new consumers, since they make the product more likely to be considered by a larger number of consumers. Thus, we argue:

Hypothesis 1: A feature update increases the number of new consumers of the product.

Reactions of existing consumers are less clear. Various evidence and arguments suggest negative reactions. To structure our discussion, we classify these arguments as rational (economic) and behavioral. "Rational" refers to the argument that actors, in terms of consumers, maximize their utility and are capable of gathering and evaluating all information for this purpose themselves [31]. "Behavioral", by contrast, denotes systematic biases in an actor's decision-making that cannot, or only in a very limited way, explained by rational logics [23]. From an economic perspective, existing consumers face costs of switching from the product they know to the updated product [17–19, 32]. If switching costs imposed by an update exceed the obtained utility, consumers may react negatively to the update. Based on Nilssen [32] and Klemperer [18], there are at least two types of switching costs relevant for our consideration of updates, namely transaction costs and learning costs. Transaction costs are immediate costs incurred by the update, in terms of the time, effort, and money in changing from the original product to its updated version [16, 18, 33]. Transaction costs include fees producers charge for the update. For example, some providers of mobile games charge for new game levels and gimmicks. Other transaction costs incur for data plans or opportunity costs associated with installing an update and fixing potential errors encountered in this process. Learning costs represent the effort required by a consumer "to reach the same level of comfort or facility with a new product as they had for an old product" [16]. For example, a feature update for a banking app might introduce a new verification procedure for making money transfers, which requires consumers to understand and learn the new procedures before being again able to transfer money. Another example might be that feature updates imply a reorganization of the user interface, eventually confronting consumers with costs for learning and understanding the handling the product.

Various studies yielded evidence for consumers' economic considerations. Several studies document learning costs in the context of product features [11, 12, 34, 35]. Thompson et al. [34] asked study participants to choose between three variants of a digital device. More than sixty per cent of participants chose the variant with the most features. Similarly, when the researchers gave subjects the chance to customize their product, freely choosing from twenty-five features, subjects also maximized features. However, when the researchers asked the subjects to use the device, subjects evaluated products with many features more negatively than the ones with less features. Mukherjee and Hoyer [12] observe in the case of high-complexity products that additional features reduced product ratings because consumers made learning-cost inferences about these features. Finally, the findings of Meyer et al. [11] suggest that while consumers are more likely to adopt products with added features, they subsequently avoid using these features due to inferred learning costs. In sum, switching costs imposed by an update might explain negative reactions of existing consumers.

Negative reactions of existing consumers may also represent a behavioral phenomenon that exists aside from rationality. If the negative reactions are behavioral, consumers react negatively even if updates provided new capabilities and implying no costs at all. Among others, behavioral drivers include psychological ownership [20, 21], reluctance toward novelty [22], loss aversion or endowment effects [24], and routine-seeking behavior [25]. These behaviors have been documented across disciplines, and they provide theoretical arguments that existing consumers evaluate feature updates negatively. Transferred to our context, existing consumers may react negatively to updates because they represent novelty in the products they use. In a similar way, endowment or loss aversion tendencies of consumers might let them prefer keeping the "status quo" of their product rather than being provided with potentially useful new features [23, 24]. In sum, we argue:

Hypothesis 2: A feature update decreases existing consumers' evaluations of the product.

3 Method

We test our hypotheses in the context of mobile apps. Apps are a type of software for a specific and particular purpose, optimized for mobile devices [36, 37]. Typical examples of apps are email, calendar, stock market, and weather. Apps optimize the appearance of displayed data, taking into consideration the screen size and resolution [36]. The functionality of mobile apps is usually limited by the unique characteristics of mobile devices: they have comparably little processing power, are controlled by touch gestures, and used "on the go" [36]. We particularly study mobile apps that run on Google's Android platform. At the time of our study, more than 80% of all smartphones worldwide run Android (Gartner, 2016). This setting has the advantage that we can collect data directly from the Google Play Store, the largest store for Android apps. In Google Play, consumers can compare, rate, review, and obtain apps [38]. Producers can update their apps at any time. Updates are rolled out immediately and automatically to existing consumers "over-the-air". Consumers are not charged any direct costs for an update, yet producers may adjust upfront prices or in-app prices along with an update.

To allow for causal inference, we employ a matching strategy [26, 28, 39]. Matching strategies pair each observation that experiences the treatment of interest at a given point in time (in our case, apps that experience an update) with one or several similar observations that do not experience the treatment at that time (the control group). We observe each app at four subsequent points in time: two weeks before the update (t-2), one week before the update (t-1), one week after the update (t+1), and two weeks after the update (t+2). We estimate the effects of updates by calculating the difference-in-differences (DID) between updated and not-updated apps, before and after the update [28, 40].

We obtained a list of all apps in the Google Play Store as of June 2016 from a mobile analytics firm. We selected a random sample of 100,000 apps from this list, for which we collected app-specific information, including ratings, updates, prices, and text reviews, in an automated way in a weekly panel format. We filtered the obtained dataset as follows. Besides apps, Google Play lists content, including television shows, music, and books, and hedonic applications, including games. In order to ensure comparability, we excluded apps labeled as "books & references", "comics", "education", "libraries & demos", "news & magazines", "wallpaper", "widgets", and “games”. To ensure comparability, we dropped apps with less than ten downloads. In the following, we discuss the variables included in our study in more detail.

DOWNLOADS. To assess whether updates attract new consumers we use the number of downloads for an app. Google Play, as other app markets, does not provide precise measures of app downloads. Instead, Google gives a categorical indicator of the number of downloads (e.g., 5-10, 100-500, 500,000-1,000,000). To obtain a more detailed measure of downloads, we combined information on download intervals with the number of ratings for an app. In order to submit a rating, users must have downloaded an app, so the number of ratings can be considered a conservative lower bound to consumer demand [41]. The resulting variable DOWNLOADS is then the mean between the midpoint of the download interval for an app and its number of ratings, which we logged.

RATING. We assess consumer reactions to updates by the rating consumers give to apps. Consumers may evaluate apps by rating it from one to five “stars”, where one star represents a low rating and five stars represent a high rating of the app. Apps with higher ratings are perceived to fulfill user expectations, have an agreeable and engaging interface, and are well-suited to audiences’ needs [38]. Consumers can renew their ratings after an app update, which allows us to distinguish between ratings of existing consumers versus new consumers [38]. The Google Play Store provides the mean rating of all consumers, rounded to one decimal, which we report as RATING.

Focal predictors (UPDATE and AFTER). To identify feature updates, we hired two independent assistants who manually inspected the changelogs (or, release notes) app producers publish along an update. In changelogs, producers describe key aspects of an app update [cf. 42]. Prior work used version numbers (e.g., 2.0, 2.1) to identify updates, which may serve as a proxy of feature updates [e.g., 5]. Although it is an informal convention that integer increases in version numbers indicate major changes [42], this standard is not enforced in many contexts and subject to certain ambiguity. Moreover, version numbers do not allow inferring the extent of features added to an app. By contrast, changelogs provide detailed insights into the changes made [42]. In the Google Play Store, changelogs are displayed below the product description in a section entitled “What’s new”, which makes them an important aspect of producers’ communication. Changelogs are limited to 500 characters, which requires producers to precisely describe the update [38].

The central predictor in our model is the dichotomous variable UPDATE, which is one if the focal app was updated with a new feature and zero otherwise. DID analyses require a second indicator for distinguishing the periods before and after the event that is studied. Thus, we include the dichotomous indicator AFTER in our models, which is one for the weeks after the update. The DID estimator is then given by interacting AFTER with UPDATE.

We construct further variables to gain insights into economic and behavioral reasoning behind negative ratings. First, the variable FEATURES ADDED is the count of the features added in an update. Second, we count the number of words in the changelog, as captured in WORDS IN CHANGELOG. Third, we obtained the time since the last feature addition. The variable WEEKS SINCE LAST UPDATE counts the number of days since the last update. Fourth, we measure direct price changes associated with an update (PRICE, continuous, in US-Dollar). The variable FREE is an indicator of apps that charge no up-front price.

Finally, consumers can include text reviews with their ratings, and these reviews may offer further insights. If reactions to updates are behaviorally driven, consumers' textual responses may give an indication. To analyze consumer reviews we implemented natural language processing techniques. We use the standard semantic text analysis software LIWC [43] to capture major text semantics. We are particularly interested in capturing consumer reasoning in a review. LIWC offers for this purpose the category "affective processes" [43]. Technically, each category consists of a list of identifying words. LIWC scores the wordlists against a text, and subsequently assigns a numerical score depending on how many of the category words were observed in a text. The category we employ in our analyses is "affective mechanism". It includes 1393 word stems, among them "happy", "worried", "hate" or "ugly". In a five-word text, "I hate the new update", the output by LIWC is 20 (per cent) for the affective mechanism dictionary (i.e., one affective word "hate" divided by a total of five words in the text, and multiplying by 100%). Before conducting the LIWC analyses, we cleansed the reviews. We removed fill words from the text, lemmatized each word, and removed reviews that were not written in English language. We then score the affective mechanism dictionary against consumers' review texts for each app-week. This procedure resulted in a numerical score assigned to each app-week, which indicates the mean score for the affective category. We include this score as the continuous variable AFFECTIVE in our analyses.

Controls. We estimate our models with app-level and time (i.e., week) fixed effects. App fixed effects adjust for static differences among apps (e.g., functionality, usability, producer etc.). Time effects control for external events (e.g., announcements by Google) or trends (e.g., an increasing number of apps are published), in terms of that app producers vary their decisions to update in response to temporal events or short-term trends.

We followed Shadish et al. [27] to build our matched set of control apps. We relied on propensity score matching [26]. We match on observational characteristics and the time of the update. The critical task in both PSM is to choose matching criteria. Matching criteria are inherently context-specific [28]. Although our context has received some attention in prior literature [e.g., 36, 37, 41], evidence is scarce when it comes to indicators of updating. We therefore followed the procedure employed by Pahnke et al. [44] and used informal interviews with producers, analysts, and industry experts to derive suitable matching criteria. The interviews converged on a number of app-specific factors. The interviews revealed that updates are costly for producers. Producers may thus tend to invest only in "promising" apps, depending on downloads and ratings received. Thus, we used DOWNLOADS and RATING as criteria for the matching procedure. Producers also seemed to update apps more often for which they charged an upfront price. Thus, we added PRICE as a matching criterion. While not evident from our interviews, we also added WEEKS SINCE LAST UPDATE as a matching criterion to account for temporal differences.

We used nearest neighbor matching without replacement in the first week of the pretreatment period. The final sample consists of 68,988 app-weeks, spanning a period of 24 weeks from 1 July 2016 to 15 December 2016. We used Ordinary Least Squares (OLS) with heteroscedasticity-robust standard errors clustered on app to estimate the following baseline equation:

$$y_{i,t} = \beta_0 + \beta_1 \mathrm{AFTER}_t \text{ x } \mathrm{UPDATE}_i + V_i + \tau_t + \epsilon_{i,t} \quad (1)$$

The subscripts i and t index for app and week, respectively. The dependent variable is $y_{i,t}$. UPDATE_i is an indicator variable for whether app i is in the treatment group, AFTER_t equals 1 if the current week is after the treatment, V_i are app fixed effects and T_t are time fixed effect. The DID coefficient of interest is β_1, which can be interpreted as the relative change of the treatment group compared to the control group, caused by the treatment. The main effects, AFTER and UPDATE, are absorbed.

4 Results

We first turn toward analyzing the consequences of feature updates for app downloads. We estimate equation (1) with DOWNLOADS as dependent variable, which gives us the effect of feature updates on app downloads. Table 1 shows the results. In Model 1, we observe that the interaction of interest, AFTER x UPDATE is positive and significant. All other things being equal, a feature update increases downloads by approximately 1.8% on average, plus minus .3%. This finding indicates that feature updates attract new consumers, supporting Hypothesis 1.

Table 1: Main Results: The Effect of Updates on App Downloads and Ratings.

	(1) Log(Downloads)	*(2) Rating*
After x Update	$.018^{***}$	$-.040^{***}$
	(.003)	(.002)
Log(Downloads)		$.062^{***}$
		(.012)
Constant	9.202^{***}	3.560^{***}
	(.009)	(.110)
Specification	OLS	OLS
Adjusted R2	.02	.05

Note: Accounts for app and time fixed effects. *, **, *** indicate significance at the 5%, 1%, and .1% levels, respectively.

How do existing consumers react to feature updates in terms of ratings? We estimate equation (1) with RATING as dependent variable, in terms of the effect of feature updates on ratings by existing and new consumers. To isolate the effect of product updates on existing consumers' ratings, we control for the increase in an app's new consumers with DOWNLOADS. In Model 2, the coefficient of AFTER x UPDATE now gives the effect of feature updates on existing consumers' ratings. The coefficient of AFTER x UPDATE is negative and strongly significant. All other things being equal, existing consumers rate an app 1.1% more negative after an update, on average. Thus, existing consumers react negatively to feature updates, supporting Hypothesis 2. So far, our findings indicate that consumers rate updated apps worse than before the update. What explains this discount? Our literature background presented economic and behavioral explanations, which we seek to explore in our data in the following. As it is infeasible to design a formal test that allows fully rejecting either economic or behavioral reasons, we conduct counterfactual analyses. We warrant, however, that our analyses do not allow definite conclusions. Rather we conduct these analyses to provide a more detailed picture of the mechanisms in place. If the negative reactions of existing consumers are a behavioral phenomenon, then the semantical analyses of consumer reviews should provide us with an indicator. To explore behavioral drivers, we assessed the semantic text analyses of consumer reviews. If updates cause an increase in the usage of affective words in reviews of existing consumers, then we should have an indicator of a behavioral bias. Figure 1 plots AFFECT for updated and not updated apps, before and after the update. The plot shows an increasing number of affective words (e.g., hate, annoyed, bad) in consumer reviews for updated apps, whereas control apps remain almost at pre-update levels. The figure indicates an increase in affective vocabulary in existing consumers' reviews for updated apps. Econometric analyses confirm this observation.

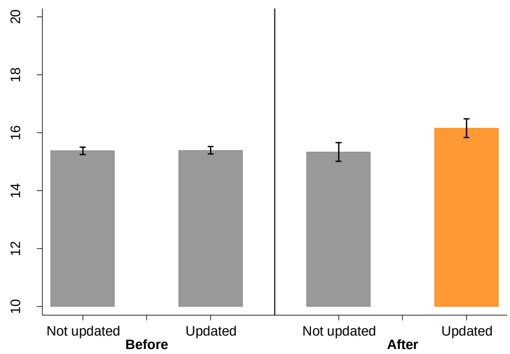

Figure 1: Mechanisms: Content of Consumer Reviews Before and After a Feature Update.

5 Discussion and Conclusion

Do updates cause greater product demand and appeal? In our data, updates appeared as an instrument to create new demand. All other things being equal, we find a feature update to increase downloads by approximately 1.8% on average. When considering the reactions of existing consumers, we found evidence of negative effects. All other things being equal, an update causes ratings of existing consumers to decline by approximately 1.1% on average, when compared with their ratings prior the update and to similar but not updated apps. These findings remained robust to various matching strategies and parameters, and account for app and time heterogeneity. In subsequent analyses, we found support that negative reactions of consumers are a behavioral phenomenon.

This paper directly responds to calls for understanding the management of digital products [1, 3, 8] and the "strategic choices related to how value is created and captured" [2]. The strategic choice we investigated is whether firms should invest in feature updates over the lifetime of their products—a phenomenon that has received increasing interest of scholars and managers. Our findings align with the observation in prior work that updates attract new consumers, yet object the observation that updates are perceived positively by existing consumers [4, 5].

Our paper sought to disentangle economic and behavioral drivers behind consumer reactions. If future research can confirm our finding of behavioral triggers behind existing consumers' negative reactions, producers' possibilities in reducing negative reactions might be limited. Here, we are left to speculate whether these reactions reflect an "update resistance", which might be a phenomenon of interest for product management [1, 3], consumer choice [e.g., 12–14], and software management literatures [e.g., 45–48]. The behavioral mechanisms actually underlying such an update resistance—such as reluctance toward novelty [22] or endowment [23]—require further decoding in future studies. Interviews with consumers and subsequent, experimental testing may help teasing out these mechanisms.

Finally, our context offered insights in the mobile app industry [37, 41, 49]. Our

Our findings suggest actionable patterns for managers. First, our findings support the effectiveness of feature updates in attracting new consumers to a product. The purposeful design of feature updates to attract particular consumer segments may allow producers to continuously adjust their products to various consumer groups as well as to be a "moving target" for competition. Second, our data do only little support economic reasons behind existing consumers' negative reactions. At least in our very context, the success of mitigation strategies focused on reducing the switching costs invoked by a feature update—such as reducing the extent and frequency of an update or providing guidelines and tutorials—might be limited.

References

1. Iyer, B., Venkatraman, N.V.: What Comes After Smart Products, https://hbr.org/2015/07/what-comes-after-smart-products.
2. Porter, M.E., Heppelmann, J.E.: How Smart, Connected Products Are Transforming Competition. Harvard Business Review. 92, 64–88 (2014).
3. Yoo, Y., Henfridsson, O., Lyytinen, K.: Research Commentary - The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research. Information Systems Research. 21, 724–735 (2010).
4. Fleischmann, M., Amirpur, M., Grupp, T., Benlian, A., Hess, T.: The Role of Software Updates in Information Systems Continuance: An Experimental Study from a User Perspective. Decision Support Systems. 83, 83–96 (2016).
5. Tiwana, A.: Evolutionary Competition in Platform Ecosystems. Information Systems Research. 26, 266–281 (2015).
6. Svahn, F., Mathiassen, L., Lindgren, R.: Embracing Digital Innovation in Incumbent Firms: How Volvo Cars Managed Competing Concerns. MIS Quarterly. (forthcoming).

7. Taub, E.A.: Your Car's New Software Is Ready. Update Now?, http://www.nytimes.com/2016/09/09/automobiles/your-cars-new-software-is-ready-update-now.html.
8. Bharadwaj, A., El Sawy, O.A., Pavlou, P.A., Venkatraman, N.: Digital Business Strategy: Toward a Next Generation of Insights. MIS Quarterly. 37, 471–482 (2013).
9. Green, P., Srinivasan, V.: Conjoint Analysis in Consumer Research: Issues and Outlook. Journal of Consumer Research. 5, 103–123 (1978).
10. Lancaster, K.: Consumer Demand: A New Approach. Columbia Universiry Press, New York, NY (1971).
11. Meyer, R.J., Zhao, S., Han, J.K.: Biases in Valuation Vs. Usage of Innovative Product Features. Marketing Science. 27, 1083–1096 (2008).
12. Mukherjee, A., Hoyer, W.D.: The Effect of Novel Attributes on Product Evaluation. Journal of Consumer Research. 28, 462–472 (2001).
13. Nowlis, S.M., Simonson, I.: The Effect of New Product Features on Brand Choice. Journal of Marketing Research. 33, 36–46 (1996).
14. Sela, A., Berger, J.: How Attribute Quantity Influences Option Choice. Journal of Marketing Research. 49, 942–953 (2012).
15. Surowiecki, J.: Feature Presentation, http://www.newyorker.com/talk/financial/2007/05/28/070528ta_ talk_surowiecki.
16. Chen, P.-Y., Hitt, L.M.: Measuring Switching Costs and the Determinants of Customer Retention in Internet-Enabled Businesses: A Study of the Online Brokerage Industry. Information Systems Research. 13, 255–274 (2002).
17. Farrell, J., Klemperer, P.: Coordination and Lock-in: Competition with Switching Costs and Network Effects. In: Handbook of Industrial Organization. pp. 1967–2072 (2007).
18. Klemperer, P.: Markets with Consumer Switching Costs. The Quarterly Journal of Economics. 102, 375–394 (1987).
19. Shapiro, C., Varian, H.R.: Information Rules: A Strategic Guide to the Network Economy. Harvard Business Press, Cambridge, MA (1998).
20. Pierce, J.L., Kostova, T., Dirks, K.T.: Toward a Theory of Psychological Ownership in Organizations. Academy of Management Review. 26, 298–310 (2001).
21. Pierce, J.L., Delbecq, A.L.: Organization Structure, Individual Attitudes and Innovation. Academy of Management Review. 2, 27–37 (1977).
22. Boudreau, K.J., Guinan, E.C., Lakhani, K.R., Riedl, C.: Looking Across and Looking Beyond the Knowledge Frontier: Intellectual Distance, Novelty, and Resource Allocation in Science. Management Science. (forthcoming).
23. Kahneman, D.: Maps of Bounded Rationality: Psychology for Behavioral Economics. The American Economic Review. 93, 1449–1475 (2003).
24. Kahneman, D., Knetsch, J.L., Thaler, R.H.: Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias. The Journal of Economic Perspectives. 5, 193–206 (1991).
25. Oreg, S.: Resistance to Change: Developing an Individual Differences Measure. Journal of Applied Psychology. 88, 680–693 (2003).
26. Rosenbaum, P.R., Rubin, D.B.: The Central Role of the Propensity Score in Observational Studies for Causal Effects. Biometrika. 70, 41–55 (1983).
27. Shadish, W.R., Cook, T.D., Campbell, D.T.: Experimental and Quasi-Experimental Designs for Generalized Causal Inference. Houghton Mifflin Company, Boston, USA (2002).
28. Angrist, J.D., Pischke, J.-S.: Mostly Harmless Econometrics: An Empiricist's Companion. Princeton Press, Princeton, NJ (2009).
29. Clements, P., Northrop, L.: Software Product Lines: Practices and Patterns. Addison-Wesley, Boston, MA (2002).
30. Ulrich, K.: The Role of Product Architecture in the Manufacturing Firm. Research Policy. 24, 419–440 (1995).

31. Kahneman, D., Tversky, A.: Prospect Theory: An Analysis of Decision Under Risk. Econometrica. 47, 263–291 (1979).
32. Nilssen, T.: Two Kinds of Consumer Switching Costs. The RAND Journal of Economics. 23, 579–589 (1992).
33. Whitten, D., Wakefield, R.L.: Measuring Switching Costs in IT Outsourcing Services. The Journal of Strategic Information Systems. 15, 219–248 (2006).
34. Thompson, D.V., Hamilton, R.W., Rust, R.T.: Feature Fatigue: When Product Capabilities Become Too Much of a Good Thing. Journal of Marketing Research. 42, 431–442 (2005).
35. Thompson, D.V., Norton, M.I.: The Social Utility of Feature Creep. Journal of Marketing Research. 48, 555–565 (2011).
36. Ghose, A., Goldfarb, A., Han, S.P.: How Is the Mobile Internet Different? Search Costs and Local Activities. Information Systems Research. 24, 613–631 (2013).
37. Ghose, A., Han, S.P.: An Empirical Analysis of User Content Generation and Usage Behavior on the Mobile Internet. Management Science. 57, 16711691 (2011).
38. Salz, P.-A.: Apponomics: An Insider's Guide to a Billion Dollar App Industry. Tyson Press, Chester, UK (2014).
39. Dehejia, R.H., Wahba, S.: Propensity Score-Matching Methods for Nonexperimental Causal Studies. Review of Economics and Statistics. 84, 151–161 (2002).
40. Bertrand, M., Duflo, E., Mullainathan, S.: How Much Should We Trust Differences-in-Differences Estimates? The Quarterly Journal of Economics. 119, 249–275 (2004).
41. Yin, P.-L., Davis, J.P., Muzyrya, Y.: Entrepreneurial Innovation: Killer Apps in the iPhone Ecosystem. American Economic Review. 104, 255–59 (2014).
42. Kemerer, C.F., Slaughter, S.: An Empirical Approach to Studying Software Evolution. IEEE Transactions on Software Engineering. 25, 493–509 (1999).
43. Pennebaker, J.W., Boyd, R.L., Jordan, K., Blackburn, K.: The Development and Psychometric Properties of LIWC2015. UT Faculty/Researcher Works. (2015).
44. Pahnke, E.C., Katila, R., Eisenhardt, K.M.: Who Takes You to the Dance? How Partners' Institutional Logics Influence Innovation in Young Firms. Administrative Science Quarterly. 1–38 (2015).
45. Banker, R.D., Davis, G.B., Slaughter, S.A.: Software Development Practices, Software Complexity, and Software Maintenance Performance: A Field Study. Management Science. 44, 433–450 (1998).
46. Barry, E.J., Kemerer, C.F., Slaughter, S.A.: Environmental Volatility, Development Decision, and Software Volatility: A Longitudinal Analysis. Management Science. 52, 448–464 (2006).
47. Kannan, K., Rahman, M.S., Tawarmalani, M.: Economic and Policy Implications of Restricted Patch Distribution. Management Science. 62, 3161–3182 (2016).
48. Telang, R., Wattal, S.: An Empirical Analysis of the Impact of Software Vulnerability Announcements on Firm Stock Price. IEEE Transactions on Software Engineering. 33, 544–557 (2007).
49. Bresnahan, T., Davis, J.P., Yin, P.-L.: Economic Value Creation in Mobile Applications. Stanford University Graduate School of Business Research Paper No. 14-08. (2014).