



**Applying Random Coefficient
Models to Strategy Research:
Testing For Firm Heterogeneity,
Predicting Firm-Specific Coefficients,
and Estimating Strategy Trade-Offs**

**Juan Alcácer
Wilbur Chung
Ashton Hawk
Gonçalo Pacheco-de-Almeida**

Working Paper

14-022

August 30, 2013

Copyright © 2013 by Juan Alcácer, Wilbur Chung, Ashton Hawk, and Gonçalo Pacheco-de-Almeida

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

APPLYING RANDOM COEFFICIENT MODELS TO STRATEGY
RESEARCH: TESTING FOR FIRM HETEROGENEITY, PREDICTING FIRM-
SPECIFIC COEFFICIENTS, AND ESTIMATING STRATEGY TRADE-OFFS

Juan Alcácer
Harvard Business School

Wilbur Chung
R.H. Smith School of Business - University of Maryland

Ashton Hawk
Fisher College of Business - Ohio State University

Gonçalo Pacheco-de-Almeida
HEC – Paris

August 30, 2013

Abstract:

Although Strategy research aims to understand how firm actions have *differential* effects on performance, most empirical research estimates the *average* effects of these actions across firms. This paper promotes Random Coefficients Models (RCMs) as an ideal empirical methodology to study firm heterogeneity in Strategy research. Specifically, we highlight and illustrate three main benefits that RCMs offer to Strategy researchers—testing firm heterogeneity, predicting firm-specific effects, and estimating trade-offs in strategy—using both synthetic and actual datasets. These examples showcase the potential uses of RCMs to test and build theory in Strategy, as well as to perform exploratory and definitive analyses of firm heterogeneity.

Authors are listed in alphabetical order; contributions were equal. For helpful comments on earlier drafts we are grateful to Rajshree Agarwal, Bill Simpson, Elon Kohlberg, Hong Luo, and seminar attendees at the Robert H. Smith School of Business (University of Maryland) and Harvard Business School.

INTRODUCTION

Textbooks generally define firm strategy as a set of decisions focused on managing organizational trade-offs in order to achieve long-term competitive advantage. Since the 1950s, theories in the academic field of Strategy have attempted to identify these trade-offs in order to understand the differential effects of firms' actions on performance (Ghemawat, 2002). Yet, empirical work in Strategy has diverged from the field's primary goal: Our models theorize why the same actions by different firms lead to different effects on firm performance, but our empirical work typically estimates the *average* effect of an action across firms. In fact, using standard regression analysis, there is no direct way to test if firm heterogeneity causes an explanatory variable's marginal effects to vary from one firm to the next. Strategy scholars have looked for alternative approaches to model firm-specific effects (e.g. firm-specific fixed effects, interaction terms between the explanatory variables and dummies for all firms, sub-sample analysis) but these are seldom ideal. This crude treatment of firm heterogeneity in marginal effects has curbed our empirical understanding of the strategic trade-offs behind our theories. What's more, it has opened a chasm between theoretical and empirical research in Strategy that threatens to undermine our accountability as business educators: If our empirical methods are ill-suited to study the effects of firm heterogeneity, how can we be sure the case studies we teach in the MBA classroom are about companies worth studying?

In this paper, we discuss how Random Coefficient Models (RCMs) can close the gap between theoretical and empirical research in Strategy. This next-generation methodology has already been used extensively in education, biostatistics, political science, and other fields to identify, model, and leverage unobserved heterogeneity at the individual level. We believe its application in Strategy would significantly advance our field by allowing scholars to directly test the many Strategy theories based on firm-level heterogeneity. This leap forward in empirical

methods should, in turn, elicit new and more granular, firm-level theories that will further our understanding of the basic principles governing competitive advantage.

In non-technical terms, RCMs can be described as generalized versions of standard methods such as OLS, probit, logit, multinomial logit, etc. It offers three important, additional features that enable Strategy scholars to (1) model firm actions that have differential effects on performance across firms, (2) identify firm-specific effects, and (3) understand the underlying strategic trade-offs that constrain firms' actions. These three features align strikingly well with the primary research goals of Strategy scholars. Specifically, the first feature of RCMs allows model coefficients to vary by firm (or group, firm cluster, agent, etc.) so that the effect of any explanatory variable on the dependent variable can differ from one firm to the next. This is a departure from standard regression models where coefficients are assumed to be constant across firms, with each constant coefficient representing the average effect of the variable for all firms in the sample. For each random coefficient in the model, RCMs generate two statistical moments: its mean, analogous to the constant coefficient in standard models, and its variance, or how much each observation deviates from the mean effect. Regarding the second key feature, RCMs can also predict firm-specific coefficients from the distribution of random coefficients, allowing researchers to identify the idiosyncratic marginal effect of any explanatory variable for a specific firm. A third key feature of RCMs allows the random coefficients in the model to be freely correlated without imposing additional structural assumptions. If two model coefficients are correlated due to a theoretical trade-off, the marginal effects of their respective explanatory variables move jointly across firms (positively or negatively). In this paper, we refer to these co-movements as coefficient covariances. These covariances among model coefficients are also a novelty *via-à-vis* standard regression analyses, where coefficients are kept constant.

In practical terms, these three features of RCMs offer several advantages over standard regression analysis that make this econometric methodology particularly suitable for empirical Strategy research. For example, using RCMs, researchers can make a critical distinction between firm actions (or explanatory variables) that are *statistically significant* and those that are *strategically significant*. The latter consist of variables that—independent of their statistical significance—exert a differential effect on performance across firms. This allows econometricians to directly test if a firm action can drive firm heterogeneity and, thus, competitive advantage. By revealing the presence of heterogeneity in the marginal effects of explanatory variables, researchers are no longer tied to the unrealistic assumption that all firm differences are perfectly captured by firm fixed effects.

The fact that RCMs also allow econometricians to directly predict firm-specific coefficients—instead of just the average effects of the explanatory variables—carves out space for a set of exciting new applications in Strategy research. By comparing the distribution of model coefficients across different firms, RCMs can help reconcile prior contradictory empirical findings grounded on the interpretation of average effects, which are sample-sensitive. In other words, RCMs may put an end to long-standing empirical debates waged over opposing average marginal effects by simply showing that marginal effects vary across firms. The estimation of firm-specific coefficient heterogeneity will also help us identify outlier firms that are worth studying in-depth using qualitative case studies. In this sense, RCMs can provide large-sample support to case-study analyses. In addition, knowing firm-specific coefficients will allow us to make more informed resource-allocation decisions. Hard evidence of the differential marginal efficiency of different firms or subsidiaries will guide CEOs of large companies and stock-market analysts as they invest scarce financial and organizational resources. Finally, using RCMs

to estimate correlations between random coefficients in a model will allow researchers to infer strategic trade-offs lurking beneath the surface of firm decisions. This feature of RCMs will shed light on the central concepts of strategic complexity and strategic fit. A better understanding of the interdependencies among firm actions will improve our knowledge of the ruggedness of business landscapes, thereby advancing the learning algorithms associated with strategic search and adaptation. Ultimately, it will lead to superior normative strategy prescriptions.

The econometric and strategic advantages of RCMs are discussed in detail in the first three sections of the paper. In the next section, we illustrate the advantages of RCM using two datasets: a set of simulated samples and a dataset based on Compustat. This empirical approach uses readily available data that facilitates replication, providing researchers with hands-on examples that ease their introduction to RCM. Our last section summarizes the paper and offers further guidelines for using RCM effectively in Strategy research.

ANALYZING HETEROGENEITY IN STRATEGY

At root, Strategy is about the causes and consequences of firm heterogeneity. Analyzing firm heterogeneity is the glue that binds Strategy research contributions together in a unified field of study. It is also what separates Strategy from the core disciplines it grew out of, and particularly from Economics. Strategy is essentially an empirical field, and it relies heavily on econometrics. But standard econometric methods are unfit to estimate or predict firm heterogeneity in the marginal effects of firm decisions, or to explain why the impact of certain decisions might vary from one firm to the next. This is because econometrics was originally engineered to estimate *average* effects across a sample of firms and to answer questions in Economics that typically revolve around firms' absolute performance or profitability. In other words, while Strategy is

concerned with the emergence of heterogeneity in firm performance—what we call competitive advantage—the field of Economics is not.

Standard treatments of firm heterogeneity in the strategy literature

The standard approach to panel data in most empirical research in Strategy and Economics is to estimate the following general model:

$$Y_{it} = \mathbf{X}_{it}\beta + \varepsilon_{it} \quad (1)$$

In equation (1), Y_{it} is the dependent variable and \mathbf{X}_{it} represents a set of K independent variables in the model. All variables have the traditional panel data structure with T_i observations (each denoted by t) for each firm i , for a total of I firms in the sample. The set of coefficients β can be interpreted as the effect of each independent variable $X_{k,it}$ in \mathbf{X}_{it} on the dependent variable Y_{it} . Thus, the total number of observations in the sample equals $N = \sum_{i=1}^I T_i$. Y_{it} is a vector ($N \times 1$), \mathbf{X}_{it} is a matrix ($N \times K$), β is a vector ($K \times 1$), and ε_{it} is a well-behaved error term vector ($N \times 1$).

It is worth noting that equation (1) represents a statistical ‘one-size-fits-all’ approach because the β coefficients in the model are assumed to be common across all firms in the sample. This assumption will be flawed in Strategy settings, in which firm heterogeneity is key. For example, prior studies have empirically documented substantial variation in the returns that different firms earn from FDI location decisions (Chung and Alcacer, 2002), investment in R&D (Knott, 2008; Knott and Posen, 2009), the choice of investment speed (Pacheco-de-Almeida, Hawk, and Yeung, 2013), and corporate diversification (Anand and Byzelov, 2012).

Strategy researchers have traditionally accommodated settings with firm heterogeneity by using shortcuts to ‘absorb’ this firm-specific variation with firm fixed effects. However, firm fixed effects are only an efficient technology to capture one specific type of firm heterogeneity:

when unobserved firm differences lead to different *levels* of the dependent variable. In this case, firm fixed effects change the regression intercepts to adjust for firm variation. When, instead, firm differences influence the *marginal effects* of a regressor on the dependent variable, firm fixed effects are an insufficient econometric response, and often constitute model misspecification. In fact, firm fixed effects regularly mask heterogeneity that is related to marginal effects rather than level effects.¹ For example, after Henderson and Cockburn noted the large effect of firm-dummy variables in their study of research productivity in the pharmaceutical industry (1996), the authors found that a deeper exploration of the drivers behind those significant firm dummies produced a more nuanced view of research productivity in pharmaceuticals (Henderson and Cockburn, 1994). This approach of unpacking heterogeneity into its component parts emerges naturally using RCMs.

RCMs and firm heterogeneity in the model coefficients

In settings where researchers hypothesize that firm heterogeneity influences the marginal effects of the model coefficients, the correct model to estimate is an RCM of the following form:

$$Y_{it} = \mathbf{X}_{it}\beta_i + \varepsilon_{it} \quad (2)$$

or, expanding the explanatory variables, of

$$Y_{it} = \beta_{1,i} X_{1,it} + \beta_{2,i} X_{2,it} + \cdots + \beta_{K,i} X_{K,it} + \varepsilon_{it} \quad (3)$$

Equations (2) and (3) are identical to equation (1) with the exception that the model coefficients β_i are now assumed to vary by subject or firm i . These model coefficients β_i are randomly distributed with mean β and variance σ , and β_i can be rewritten as $\beta_i = \beta + u_i$, where u_i represents the deviation of the coefficient associated with subject i from the average effect β (the common mean coefficient across all firms).

¹ Similar arguments can be made about using random effects models to capture individual heterogeneity.

Figure 1 graphically illustrates the RCM approach. Traditional methods would estimate the average effect β . RCMs assume that β comes from a distribution with mean β and variance σ and that every firm has its specific effect β_i . This firm-specific effect is the mean effect calculated from the multiple observations of a firm, Y_{it} and X_{it} (β_i as similar to the coefficient of X on Y calculated for a subsample that has only T_i observations for firm i).

Note that the simplest example of an RCM is one where only the constant in the model—and none of the other model coefficients—is random. This is equivalent to the well-known random effects regression model.

RCMs estimate two parameters for each coefficient in equation (3): (1) its mean, representing the average effect common to all firms in the sample (similar to the usual standard coefficients estimated in conventional regression models); and (2) its variance, which captures the underlying distribution of subject-specific or firm-specific coefficients. This estimated variance offers evidence of coefficient heterogeneity among firms. Thus, instead of only estimating a fixed coefficient corresponding to the first moment of the coefficient distribution (its mean), RCMs also estimate the second moment of the coefficient distribution (its variance). Both inform researchers about the distribution of the marginal effects in the model. Importantly, note that in RCMs, coefficient heterogeneity among firms is tested for rather than imposed; whether a specific model coefficient varies by firm is determined by the data rather than a researcher's pre-specification. This is because RCMs' primary objective is to assist researchers precisely when there is unobserved firm heterogeneity, that is, when econometricians intuit differences among firms but cannot identify or account for those differences in their models. In sum, RCMs are an effective way of coping with unobserved firm heterogeneity without imposing overtaxing and unwarranted model pre-specifications or estimation assumptions.

But if RCMs' claim to fame is mostly the estimation of one additional parameter—the variance of the coefficient distribution—what insights can they really bring to the field of Strategy? The answer is that the extra parameter estimated, the variance, allows Strategy researchers to make a critical distinction between explanatory variables that are *statistically significant* (e.g. their mean is different from zero) and those that are *strategically significant* (e.g. their variance is different from zero). An explanatory variable with non-zero variance reveals the presence of heterogeneity in its marginal effect: a differential effect that may be attached to a source of competitive advantage.

To better illustrate the econometric and conceptual advantages derived from estimating variances of effects, let's explore three cases.

Case 1: Mean is statistically insignificant, but variance is different from zero

This is a classic example of type II error for Strategy scholars. An explanatory variable with no average effect would be labeled unimportant even when its effect varies by firm and indicates an important driver of performance variation. Figure 2 depicts this hypothetical scenario, which is likely to occur when half of the population in the sample experiences a positive effect and half experiences a negative effect, or when the firm-specific effects are close but different from zero (as with the highly unpredictable outcomes of R&D investment, for example). We provide additional examples in our empirical section.

Note that, although the average effect is close to zero, the explanatory variable is still a source of interest in Strategy research. In this case, we say that the coefficient is not statistically significant, but it is strategically significant. Thus, adjusting our concept of significance—and, by association, our ideas about what is worth studying (or not) in Strategy— to include evidence based on variance estimates would enrich our understanding of firm heterogeneity. Imagine all

the papers that were never submitted to Strategy journals for lack of significant effects; some of them would have variances that are distinct from zero and thereby deserve further consideration for their potential impacts on firm performance. In fact, Strategy theories that have been abandoned because of failure to reject the null hypothesis of an average effect may be revived when firm heterogeneity is modeled, identified, and tested properly—that is, when both the first and second moments of the model coefficient distributions distribution are considered. For example, the empirical literature on strategic groups in the mid 90s failed to find a link between strategic group membership and firm profitability (Cool and Dierickx, 1993). As a result, interest faded in this stream of work. However, it is possible—it has in fact been hypothesized (Mas-Ruiz and Ruiz-Moreno, 2011)—that unobservable heterogeneity between strategic groups has tainted the empirical results. Specifically, strategic group membership boosts firm profits for some strategic groups, but lowers it for others. Thus, the net effect when looking across samples is often inconclusive.

Case 2: Mean is statistically significant, but variance is zero

This is a classic example of type I error for Strategy scholars. An explanatory variable with a statistically significant effect that does not vary by firm can be perceived as an important driver of performance, even when its effect is the same for all firms. Figure 3 depicts this kind of situation. We present some examples in the empirical session.

Let us emphasize that variables falling under this case are still important for firm performance and should be studied by Strategy researchers. However, if there is no difference in the effect across firms, the strategic action behind the variable may not provide a distinctive source of performance for firms.

We believe this case may become more common in Strategy as researchers get access to larger datasets. Large datasets have the obvious effect of increasing the statistical significance of

the mean coefficients, potentially increasing false positives. For example, researchers that work with patent data are used to finding high levels of significance for most variables used. By using RCMs and estimating the variance of the effect of explanatory variables, they may be able to isolate those factors that really drive their results.

At a conceptual level, some Strategy theories may have been given unduly normative credence because their claims were wrongly supported by false positives. New research using RCMs as an econometric tool, and the redefinition of what statistical significance means when two statistical moments are estimated, may provide a stricter test for normative predictions.

Case 3: Mean is statistically significant, and variance is different from zero

This case, which allows for any mean effect different from zero, is the generalization of the first case. As in case 1, the main issue here is that inferences rooted in the interpretation of average effects may not apply to all firms. Moreover, if variance among coefficients is large, findings are likely to be sample-sensitive, creating contradictory empirical findings across papers.

Examples of this problem abound in Strategy. A case-in-point is the mixed empirical findings for first-mover advantages, which have been attributed to the fact that the net benefits of being early to a market vary by firm. This observation implies firm heterogeneity in the marginal effect of entry timing on firm performance. However, recent empirical work has only tackled this issue by including moderating effects of third variables as standard interaction terms in regular regression models (for a review, see Franco *et al.*, 2009). This approach assumes that all firm attributes responsible for the heterogeneous effect of timing on performance are observable and impact performance exactly as specified by the interaction terms. Both assumptions are typically unverifiable and open to biases. Using RCMs would solve these problems by directly estimating coefficients for both mean and variance that capture each firm's net benefits from earlier market entry.

Another example central to Strategy is the effects of diversification on firm performance. A vast set of papers have offered evidence supporting both a diversification discount and a diversification premium. Anand and Byzalov (2012) used an RCM to further explore this long stream of contradictory evidence and found that the effect of deviation from the mean is roughly twice as large as the baseline diversification discount; in fact, the diversification discount becomes a premium discount for some firms. Similar findings have been produced for foreign direct investment decisions (Chung and Alcacer, 2002) and the choice of investment speed (Pacheco-de-Almeida, Hawk, and Yeung, 2013).

To conclude, embracing RCMs could spark a major change in the way Strategy researchers think about unobserved sources of firm heterogeneity and their effects on strategy fundamentals: from shifts in the model intercepts given by standard firm fixed effects to shifts in (some of) the model coefficients. For example, treating firm heterogeneity through fixed effects in a model where the dependent variable is performance is equivalent to saying, 'Firms differ because they are different.' It simply pushes back in the chain of causality, without addressing the true reason why firm performance heterogeneity, or competitive advantage, emerges. In contrast, treating firm heterogeneity through shifts in the model coefficients allows researchers the opportunity to identify competitive advantages at their origins. That is, RCMs allow the field of Strategy to study situations in which firms start out as relatively identical, but then grow differently due not only to distinct strategies but also to the fact that similar firm decisions lead to different performance outcomes. A case-in-point is Walmart. Walmart is known for squeezing margins out of suppliers because of its large market size. But this path to profitability only emerged *after* Walmart became a large company. It doesn't provide any insight into what strategies allowed

Walmart to reach a critical level of market size initially. RCM analysis of historic data when Walmart was nascent may pinpoint the timing and source of this divergence.

Finally, RCMs are more consistent with the general idea of firms as proactive strategizing and optimizing agents. Consider a study of firm productivity. It is very likely that firms that are more productive in R&D, for instance, re-optimize and invest more in R&D on the margin. RCMs are able to capture these differences in marginal effects of the regressors, while most conventional regression methods are not. Since conventional regression methods only estimate average coefficients across firms, a simultaneity bias may exist.²

PREDICTING FIRM-SPECIFIC EFFECTS

The previous section discussed how RCMs estimate both the first moment (mean) and the second moment (variance) of a model coefficient distribution. This section is motivated by a simple follow up observation: that when the second moment of a distribution is large, its first moment becomes less informative about the underlying sample. For example, with no variance, the mean completely describes the sample, but when variance is extremely large, the mean may be drastically less representative of specific sample observations. Thus, inferences based on the mean model coefficients are less meaningful when coefficients have large variance. In these cases, RCMs can do something else that conventional regression methods cannot: estimate *firm-specific* model coefficients.

Measuring firm-specific coefficients is particularly relevant in Strategy research. If coefficients vary by firm, the same action taken by different firms leads to different outcomes, thereby increasing firm heterogeneity. Some firms will be better off than others, but which firms, and how much better off? Answering these two questions is at the crux of what Strategy is about:

² See Levinsohn and Petrin (2003) for a more detailed explanation.

the ability to issue normative recommendations about how specific firms may achieve competitive advantage.

Equations (2) and (3) in the previous section are a general representation of RCMs with firm-specific coefficients. Although not providing exact firm-specific estimates, RCMs do predict firm-specific values conditional on the overall distribution of the sample estimates and on the actual data.³ The intuition of the method is captured in Figure 1 for coefficient β and firm i . The starting point is that the RCM estimates β 's overall distribution for the sample (the mean and standard deviation). Then, for each explanatory variable of strategic interest, RCMs provide not only point estimates for each firm i 's coefficient β_i (using the probability of obtaining such coefficient conditional on Y_{it} and X_{it} in an application of Bayes Theorem), but also confidence intervals around each point estimate.⁴ This is a noteworthy departure from most conventional regression methods that typically only estimate the *average* effects of strategy decisions across firms.

Firm-specific coefficients in conventional regression models versus RCMs

Conventional regression methodologies do not lend themselves to estimating firm-specific coefficients in a direct way. Indirect estimation is possible, but difficult. We outline the two indirect options below.

The first approach involves explicitly building the firm-specific coefficients into the structure of the functional form to be estimated. This essentially requires adding to equation (1) interaction terms between all the explanatory variables of strategic interest (with significant standard deviation) and firm dummies, a task that is unfeasible or very taxing on the properties of the estimators. There are usually not enough degrees of freedom to estimate the model because of

³ For a comprehensive explanation of how individual (firm) coefficients are predicted in RCMs, see Section 15.10 from Greene's book, *Econometric Analysis*, seventh edition, pp. 642–650.

⁴ See Greene, 2007, pp. 644 for a theoretical explanation on how confidence intervals are calculated.

the large number of interaction terms needed to have firm-specific coefficients. When estimation is possible, obvious efficiency concerns remain.

The second empirical strategy to indirectly obtain firm-specific coefficients in conventional regression models consists of splitting the sample by firm so that each subsample only contains observations for a single firm, and then estimating equation (1) on each of the separate subsamples. This method should be essentially equivalent to the first approach, including having several disadvantages. By splitting the sample by firm we lose efficiency in the estimation of the non-random coefficients, that is, those coefficients that do not vary by firm and that should have been estimated in the pooled sample. The different regressions may also have problems in terms of the degrees of freedom available for each subsample. Finally, splitting the sample makes comparisons across subsamples harder to establish since tests require nested models.

RCMs have clear econometric advantages over the two indirect methods of estimating firm-specific coefficients using conventional regression techniques. RCMs ‘get the job done’ with just a few additional parameters, which increases estimation efficiency without taxing the non-random model coefficients. RCMs are also more lenient in response to researchers’ ignorance about which model coefficients are random: RCMs allow and test for coefficient heterogeneity, while conventional regression techniques make assumptions about coefficient heterogeneity and build them into the model structure or estimation methodology.

From a Strategy perspective, the seamless ability of RCMs to estimate firm-specific coefficients opens the door for substantive advances in research. It allows researchers to develop and test more granular, firm-level theories and hypothesize about specific firm actions that have different effects on performance for different firms. Firm-specific coefficient estimation may also provide an alternative way to cluster firms. For example, in some circumstances it may be

more sensible to group firms according to their levels of R&D productivity (i.e. analogous R&D coefficient size) rather than by how much they invest in R&D. Finally, by estimating firm-specific coefficients, RCMs can identify firms that are outliers in marginal effects, which then become natural candidates for in-depth qualitative case studies. In other words, RCMs use large sample regressions to support case study analyses. Evidence coming from this observed reality will, in turn, further increase confidence in our econometric results. By closing the research gap between large sample studies and qualitative research, RCMs will also help determine which companies are worth discussing in the context of MBA classroom teaching.

IDENTIFYING STRATEGY TRADE-OFFS IN EMPIRICAL RESEARCH

Another central goal in Strategy is to provide insights into the trade-offs behind firm decisions. Of special interest to Strategy scholars is how these trade-offs affect firm performance. RCMs are particularly well qualified to tackle this problem. In this section we discuss how analyzing covariances between random coefficients offers a novel way to study a different and important type of strategy interdependency.

Besides RCMs, researchers can follow two main empirical approaches to Strategy trade-offs. These three approaches complement, rather than substitute for, one another. Each method examines a different type of trade-off in Strategy.

Method 1: Pearson (product-moment) correlation coefficient (PCC)

The first well-known conventional method to study strategy trade-offs uses PCC to make first-order inferences about whether two variables exhibit linear correlation or probabilistic dependence:

$$\rho_{X_{1,it},X_{2,it}} = \frac{cov(X_{1,it},X_{2,it})}{\sigma_{X_{1,it}} \sigma_{X_{2,it}}} \quad (4)$$

The PCC shows whether and how the values or levels for the independent variables co-vary (e.g. over time and across firms).⁵ The value range for the correlation is $[-1,1]$. A PCC value of 0 shows no correlation, whereas extreme values of -1 and 1 indicate perfect correlation (negative and positive, respectively). Perfect correlation means that a linear equation completely describes the relationship between the two variables. The PCC is scale-invariant: any linear transformation of the two variables will not change the correlation coefficient.

Consider a study on the effect of generic strategies on firms' competitive advantage (Porter, 1980). The specific strategy trade-off we are interested in is the trade-off between differentiation and cost strategies. The dependent variable in this study is a measure of firms' competitive advantage (e.g. deviations of firm-specific returns from the industry average). Independent variables include measures of product differentiation, cost measures, and controls. A long-standing hypothesis in the literature is that most firms are not able to pursue differentiation and cost strategies simultaneously. This is related to the well-known empirical fact that the *efficiency frontier* is typically upward sloping (see, for example, Besanko, 2007).⁶ In this example, the PCC summary statistic should display negative correlation between our measures of differentiation and cost strategies. While useful, the information conveyed by the PCC is only a first step toward understanding the differentiation–cost trade-off. In particular, the PCC does not shed light on how the trade-off between our two explanatory variables affects firm performance, which is the ultimate purpose of the study in this example.

Several statistical problems with the PCC have also been raised in the literature. Interpreting the absolute value of the coefficient is often somewhat arbitrary. For example, a correlation of

⁵ Equation (4) is the standard expression for the PCC applied to a population. The PCC for a specific sample (usually denoted by r) is obtained by replacing the numerator and denominator in equation (5) by the sample estimates of the covariances and standard deviations.

⁶ An industry's efficiency frontier is defined as the lowest level of cost that is attainable to achieve a given level of differentiation with the available technology.

0.8 may be considered high or low depending on the empirical context and measurement noise. The PCC is also not as good a measure of statistical dependence when variables have nonlinear relationships. Finally, the correlation coefficient may be particularly sensitive to the sampling methodology adopted in a study.

Method 2: Interaction terms in standard regression analysis

Interaction terms in regression analysis have a long history in Strategy research. For ease of exposition, consider the simplest case of a bi-variate version of equation (1):

$$Y_{it} = \beta_1 X_{1,it} + \beta_2 X_{2,it} + \gamma X_{1,it} X_{2,it} + \varepsilon_{it} \quad (5)$$

In equation (5), the interaction between our two independent variables is represented by the pairwise-product between them. Note that, although we will focus on the simplest case of pairwise interactions, the main points in this subsection are also applicable to higher-order interactions (i.e. higher-order products).

In an interaction term, the *marginal effect* of one independent variable on the dependent variable is a function of the *value*, that is, the *level* of (or the data for) another independent variable. For example, the marginal effect of $X_{1,it}$ on Y_{it} is a linear function of the values of $X_{2,it}$ (and vice-versa): $\frac{\partial Y_{it}}{\partial X_{1,it}} = \beta_1 + \gamma X_{2,it}$. An interaction term is also the mathematical expression for what is commonly called in the Strategy literature a moderating effect, which is typically represented by the solid line in Figure 4.

In Strategy and Economics, interaction terms determine whether two variables are strategic complements or strategic substitutes (Bensanko et. al, 2007 pp. 243-245). In particular, this is defined by the sign of the cross-partial derivative between two independent variables of interest.

In equation (5), the cross-partial derivatives are identical and equal to $\frac{\partial^2 Y_{it}}{\partial X_{1,it} \partial X_{2,it}} = \frac{\partial^2 Y_{it}}{\partial X_{2,it} \partial X_{1,it}} =$

γ . When $\gamma > 0$, variables $X_{1,it}$ and $X_{2,it}$ are strategic complements because the marginal effect of

one variable is increasing in the value of the other variable. This is also called supermodularity (Topkis, 1998). The opposite is true if $\gamma < 0$, in which case the independent variables are strategic substitutes and their relationship is submodular.

Our example about a hypothetical study of the effect of generic strategies on firms' competitive advantage illustrates well the importance of interaction terms in Strategy. The trade-off between differentiation and cost strategies, which the Pearson Correlation Coefficient captured a first glimpse of, is also characterized by a more complex relationship between our differentiation and cost measures. Specifically, the Strategy literature has long hypothesized that firms that try to simultaneously pursue differentiation and cost lose focus and, thus, have lower performance. In other words, differentiation and cost leadership are typically *strategic substitutes*. This effect of the trade-off on performance is also known as being 'stuck-in-the-middle' (Porter, 1980). Statistically, this feature of the differentiation–cost trade-off is represented by an interaction term between both independent variables, where the interaction term coefficient is expected to be negative. The *marginal* returns to further pursuing cost leadership strategies should be lower for firms with higher *levels* of differentiation.

Although interaction terms are an extremely powerful method of capturing a fundamental type of Strategy trade-off, they cannot accommodate all types of interdependencies in our studies. Importantly, interaction terms are unfit to represent situations in which unobserved firm heterogeneity moderates the effect of a trade-off on the dependent variable. For example, firm capabilities are a popular intangible construct that typically moderates the effect of Strategy observables on performance in many of our Strategy theories. Yet, we either cannot measure capabilities or only have loose empirical proxies for them. And, when we do, our knowledge of how exactly these intangibles impact our dependent variables remains elusive. Therefore, any

structured specification of these effects using standard interaction terms is prone to econometric misspecification.

In sum, the use of interaction terms in standard regression analysis is a useful method to model one specific type of strategy trade-off. However, it assumes that all heterogeneity can be (a) measured and (b) delineated in the specific functional form of pairwise- (or higher-order) products between variables. This is often not the case in Strategy. In addition, standard interaction terms can be complicated to interpret, especially in non-linear models or when the interactions are defined as higher-order products (Hoetker, 2007; Zelner, 2009). Finally, the inclusion of multiple interaction terms in regression analyses frequently creates multicollinearity problems.

Method 3: Covariances between random coefficients

RCMs in equations (2) and (3) have two important advantages over most other regression methodologies. First, RCMs do not force the econometrician to be explicit about the functional form through which firm heterogeneity moderates the effects of the independent variables in the model. Second, RCMs do not even require the econometrician to identify the drivers of firm heterogeneity. This is because the model coefficients are allowed to vary randomly by firm.

This agnostic and flexible regression methodology naturally lends itself to the study of unobserved firm heterogeneity in Strategy research. As per their properties, RCMs can accommodate two main types of ‘econometric ignorance’: (a) theoretical ignorance about the exact sources of firm heterogeneity and how they are expected to impact the structure of the econometric model; and (b) empirical ignorance, or when the sources of firm heterogeneity are known but unobservable.

The influence of unobserved firm heterogeneity on the marginal effects of the independent variables in the model is captured by the covariances between the random coefficients in RCMs.

In a simple bi-variate version of RCMs, this is given by:

$$\text{cov}(\beta_{1,i}, \beta_{2,i}) \quad (6)$$

Equation (6) means that the marginal effects of the independent variables in the model co-vary. RCM estimation allows for coefficient co-variation based on the existence of latent differences between firms that simultaneously drive the partial effects of more than one independent variable. In other words, if the same unobserved differences between firms simultaneously influence the size of multiple coefficients in a systematic pattern, we should expect coefficient co-variation.

To show coefficient covariance in context, let us go back to the Strategy example about the differentiation–cost trade-off and its effect on firm performance. A sensible hypothesis is that more innovative firms benefit more from investing in differentiation or cost strategies. For firms that are better at product innovation (differentiation) or process innovation (cost reduction), each extra dollar invested in either strategy should create more value (e.g. Besanko, 2007). In Strategy terms, this means that the differentiation–cost trade-off is less binding. In econometric terms, this means that the marginal effects of differentiation and cost strategies on firm performance should positively co-vary, and that this positive coefficient co-variation is a function of unobserved firm heterogeneity in innovation capabilities. Note that coefficient covariance cannot be empirically documented by looking at simple correlations between the independent variables or standard interaction terms (methods 1 and 2 above). To recapitulate, in standard interaction terms the

marginal effect of one independent variable is a function of the *value* of another independent variable. Instead, coefficient covariance is about the correlation between *marginal* effects.⁷

In short, RCMs are a versatile econometric tool. They can easily accommodate econometricians' ignorance about how unobserved firm heterogeneity affects fundamental strategic trade-offs. RCMs are flexible to the point that coefficient covariance is allowed and tested for, but not required or assumed. Results are not pre-built in the model specification; they emerge from the data. This methodology opens up new and promising avenues for theorization about trade-offs between the *effects* of firms' decisions and not just trade-offs between the decisions themselves, as illustrated by the differentiation–cost trade-off discussed above. Finally, since RCMs naturally lend themselves to tackling the effects of unobserved firm heterogeneity, they are an excellent methodology to get at the 'deep parameters' of Strategy research. These deep parameters include many of the intangible constructs that pervade our theories (capabilities, reputation, etc.), but for which we have collectively failed to provide credible empirical measures.

APPLYING RCMs TO STRATEGY: TWO BASIC EXAMPLES

We offer two basic examples to illustrate the advantages and limitations of RCMs for Strategy scholars. First we estimated four models (cross-sectional OLS, panel data fixed effects, panel data random effects, and random coefficients) on 100 randomly-generated datasets. We refer to these datasets as synthetic datasets. Second, we estimated basic productivity functions for two industries (pharmaceuticals and automobiles) using Compustat as a data source.⁸

⁷ In method 2 with standard interaction terms, the marginal effects of two independent variables only co-vary if there is correlation between the independent variables themselves. For example, in the two-variate model with interaction terms above, $cov\left(\frac{\partial y_{it}}{\partial X_{1,it}}, \frac{\partial y_{it}}{\partial X_{2,it}}\right) = cov(\beta_1 + \gamma X_{2,it}, \beta_2 + \gamma X_{1,it}) = f(cov(X_{1,it}, X_{2,it}))$.

⁸ Synthetic datasets and software code used in this paper are available from the authors upon request.

RCMs estimated on synthetic samples

The objective of this simulation exercise is to generate samples that illustrate the three features that RCMs provide to Strategy scholars: (1) modeling firm heterogeneity, (2) predicting firm-specific coefficients, and (3) modeling strategic trade-offs through covariances of effects. We generated 100 random samples where we could estimate the following equation:

$$Y_{it} = \beta_1 X_{1,it} + \beta_{2,i} X_{2,it} + \beta_{3,i} X_{3,it} + \varepsilon_{it} \quad (7)$$

where i indices firms and t indices observations for a given firm i . In the synthetic data, $i=1..100$ and $t=1..10$ (equivalent to a dataset of 100 firms for 10 years); $\beta_1, \beta_{2,i}, \beta_{3,i}$ are firm-specific random coefficients drawn from normal distributions; $X_{1,it}, X_{2,it}, X_{3,it}$ are random explanatory variables, observed by the econometrician; and ε_{it} is an iid error term drawn from a normal distribution with mean 0 and standard deviation 1 (e.g. $\varepsilon_{it} \sim N(0,1)$).

The process of drawing our samples is described in detail in Appendix A, but some specific decisions made while building the dataset require further explanation here. First, each coefficient in equation (7) was designed to follow the cases described in the section titled “Analyzing heterogeneity in Strategy.” Specifically, the goal of β_1 is to capture the case when the effect of $X_{1,it}$ on Y_{it} is common across firms: a case of statistical significance but of no strategic significance. Therefore we set $\beta_1 \sim N(0.5,0)$. $\beta_{2,i}$ corresponds to the case when the average effect $X_{2,it}$ on Y_{it} is zero, which will translate into a lack of statistical significance under traditional estimation techniques, even when the variable is strategically significant. Thus we set $\beta_{2,i} \sim N(0,0.25)$ guaranteeing that for some firms $\beta_{2,i}$ would be positive while for others it would be negative, in both cases affecting firm performance. Finally, $\beta_{3,i} \sim N(-5,25)$, matching the case when there is a clear negative average effect (mean equal to -5), but where that effect varies widely by firms (that is, where there is a large variance). Although traditional estimation

techniques would be more or less precise at estimating the average effect of $X_{3,it}$ on Y_{it} (first moment of β_3 's distribution), they will miss the fact that, for some firms, the effect will be quite different (second moment of β_3 's distribution). Note that $X_{3,it}$ would be the case of a variable being both statistically significant and strategically significant.

Second, we aim to capture through the generation of explanatory variables the idea of firm-specific capabilities. To accomplish this goal, we drew random values of $X_{1,it}, X_{2,it}, X_{3,it}$ *only* for the first observation by firm i (e.g. $X_{1,i1}, X_{2,i1}, X_{3,i1}$) from equally normal distributions with mean 50 and standard deviation 10 (e.g. $X_{1,it} \sim N(50,100)$, $X_{2,it} \sim N(50,100)$, and $X_{3,it} \sim N(50,100)$). We populated the remaining values (e.g. from $t=2$ to $t=10$) following a random walk process from the initial values. In other words, $X_{1,it} = X_{1,it-1} + u_{it}$ with $u_{it} \sim N(0,16)$, $X_{2,it} = X_{2,it-1} + v_{it}$ with $v_{it} \sim N(0,16)$ and $X_{3,it} = X_{3,it-1} + w_{it}$ with $w_{it} \sim N(0,16)$.⁹ We use the same procedure for each of the 100 randomly-generated datasets. Table 1.a shows the descriptive statistics for all the values simulated, including explanatory variables, coefficients, error term, and dependent variable. Table 1.b shows the correlation table.

Table 2 shows the summary results from estimating a variety of models on the synthetic samples. Column 1 in Table 2 shows the true parameters of the coefficients, their standard deviations, and covariances that were used for the synthetic datasets. The table also shows the average results from 100 samples of estimating equation (7) assuming the following data structures: cross-sectional OLS (column 2), panel data with fixed effects (column 4), panel data with random effects (column 6), and random coefficient models (columns 8 and 10). Note that methods in columns 2, 4, and 6 provide average estimates (first moment of coefficient

⁹ We explored alternative ways to obtain firm-specific capabilities, such as firm-specific constants. However we decided against this option because it would imply a more complex model where we would have to specify covariance between the firm-specific constants and the remaining variables.

distributions) and control for firm heterogeneity; the RCMs in columns 8 and 10 provide average estimates as well as estimates of their standard deviations (second moment of coefficient distributions) and correlations.¹⁰

Although we present averages of both coefficients estimated and their standard deviations to construct a table with a format familiar to researchers, reporting statistical significance across the 100 samples in a meaningful way is challenging. To solve this problem, we looked at the results of each sample and report the frequency with which a specific estimate was significant across synthetic samples. Thus, column 3 represents the frequency with which estimates were significant at the 5 percent level for OLS, column 5 for panel data fixed effects, column 7 for panel data random effects, and columns 9 and 11 for RCMs. We believe this unorthodox approach provides an important element to compare models.

Inferences for estimates of variances and covariances require further comment. Performing traditional tests that explore whether estimated coefficients are equal to zero is challenging with RCMs for two reasons. First, variances always take positive values and are lower-bounded by zero. Second, if the variance for any coefficient is actually zero, it imposes extra restrictions on the correlation of that coefficient with the others.

For inferences concerning the covariance parameters, researchers can use statistics such as the Wald Z, which is computed as the parameter estimate divided by its asymptotic standard error.¹¹ Although the Wald Z may be valid for large samples, it can be unreliable for small data sets and for variance estimates, which require the possibility of negative variance components.

¹⁰ We estimated our RCMs using the mixed command in Stata 13 (xtmixed in previous versions). Stata's standard output provides standard deviations and correlations instead of variances and covariances. Other software packages that also estimate RCMs vary in their standard output. SAS' output (procedure mixed) provides variances and covariances. Limdep's output (command regre) provides standard deviations and covariances.

¹¹ The asymptotic standard errors are computed from the inverse of the Fisher information matrix, which is calculated as the expectation of the square of the derivatives of the log-likelihood (with the derivatives taken with respect to each of the covariance parameters).

Nonetheless, this is the approach we follow because its use is common in Economics (Train, 1998; Revelt and Train, 1998; Greene, 2004; Craig, Greene, and Douglas, 2003; Nevo, 2000) and Strategy (Chung and Alcacer, 2002; Pacheco-de-Almeida, Hawk, and Yeung, 2013; Knott, 2008, 2012; Knott and Posen, 2009).¹² Therefore, columns 9 and 10 indicate the percentage of models where the Wald test for the variance (standard deviation) and covariance (correlation) parameters were statistically significant.

Alternatively, researchers may follow a test based on log-likelihood ratio that compares two models: a reduced model (without variances or covariance to be estimated) and an RCM (with variances and covariance to be estimated). As long as the reduced model does not occur on the boundary of the covariance parameter space, the statistic computed follows asymptotically a χ^2 distribution with degrees of freedom equal to the difference in the number of parameters estimated between both models. If the reduced model does occur on the boundary of the covariance parameter space (e.g. the variance is equal to zero) the asymptotic distribution becomes a mixture of χ^2 distributions, whose “probability tails are bounded by those of the χ^2 distribution with degrees of freedom equal to the full number of restricted parameters” (Stata, 2013: 12). Therefore the χ^2 test represents a conservative comparison across models.

Additionally, researchers can also calculate 95 percent confidence intervals for a given parameter. A very tight interval around zero is a good indicator that the estimated variance is equal to zero. For example, the average 95 percent interval for β_1 across samples for column 8 is [0.002538, 0.060], suggesting zero-variance of the effect of X_1 on Y .¹³

¹² Limdep is an example of statistical package that also follows the norm.

¹³ Note that that these 95 percent confidence intervals around the point estimate cannot be interpreted as traditional hypothesis testing because they are not set around the null hypothesis that the coefficient equals 0. Instead they are centered on the estimated value of beta.

Returning to Table 2, note that the models based on random coefficients offer the best fit for the synthetic datasets. Based on the decrease in the average log likelihood, model 8 is better than the fixed effect model in column 4, which was the best of the traditional models, and that the RCM with covariance in column 10 is even better. In some cases, for example a cross-sectional OLS, the model is obviously a misspecification of the data structure we generated and, as such, generates the worst fit.

Model misspecification is revealed not only through measures of model fitness but also in the results for each independent variable. For example, the average estimated coefficient for X_1 (β_1) is 0.33, a value far from the true parameter of 0.5. Moreover, OLS models fail to recognize that β_1 is different from zero in 51 percent of the samples. In contrast, panel data models based on fixed effects and random effects offer better average estimates for β_1 : 0.501 and 0.500, practically identical to the true parameter. However, even when these models are closer to the true parameter on average, they fail to recognize that the sample parameter is different from zero in 30 percent and 29 percent of the samples, respectively. Surprisingly, RCMs are able to recognize that β_1 is different from zero for every single sample while also providing average estimates very close to the true parameter 0.5. Why? RCMs take full advantage of firm heterogeneity while also estimating the variance of β_1 , factoring in that it is very close to zero (coefficient of 0.008, with very small average standard deviation and 95 percent confidence intervals). Even more, in 81 percent of the samples we cannot reject the null hypothesis that the variance is equal to zero. Without knowing that there is no variance of the effect of X_1 on Y , traditional models are more likely to misallocate the underlying heterogeneity in equation (7) and fail to reject that β_1 is actually different from zero.

Figure 5 reinforces this intuition. It shows the distribution (kernel density) of the 100 estimates of β_1 obtained for each method in Table 2: OLS, and panel data using fixed effects (FE), random effects (RE), and RCM. As you can see, the distribution of β_1 estimates is really wide for OLS, with a good proportion of estimates around the value of zero. In these cases, OLS would fail to find significance for β_1 . The distributions for FE and RE are almost identical to each other and narrower than OLS'. Although an improvement, around 30 percent of the samples produce FE and RE estimates that are close to zero, triggering lacks of significance for X_I . In contrast, RCMs produce a very narrow distribution that is centered on the true parameter value, 0.5. As a result, no estimates are close to zero, guaranteeing significance in each sample.

From the above analysis we note that, even if the econometrician assumed heterogeneity for β_1 in equation 7, the results in column 8 suggest there is none (variance practically equal to zero) and that the β_1 coefficient should be fixed (β_1) rather than random ($\beta_{1,i}$). As result, we do not allow β_1 to vary by firm in the model in column (10) (nor do we introduce covariances of β_1 with β_2 or β_3 since β_1 is constant across firms). This is a case of a variable that is statistically significant but of little strategic interest: the effect of increases of X_I on Y is the same for all firms. In other words, any difference in performance is explained by differences in levels of the independent variable X_I and the effect of changes in levels is the same for all firms.

Comparing results for β_2 across models showcases an example of effects that would not be recognized as worthy of study by Strategy scholars—even when they are. It captures the case of a variable, X_2 , whose effect varies by firm even when the average effect is zero (and therefore not statistically significant at the sample level). Note that the most accurate method should estimate a sample average effect of zero and fail to reject the null hypothesis (find no significance) in as many samples as possible. OLS again is the worst performing model, with an

average coefficient across samples of 0.139 while rejecting the null hypothesis of zero effect in almost 60 percent of the samples. Panel data models with fixed and random effects perform better, with average coefficients closer to zero (0.016 and 0.017, respectively) and a failure to reject the null in two-thirds of the samples. RCMs perform the best, failing to reject the null in only 7 percent of the samples and providing average estimates close to zero. Again we capture these differences through a graph (Figure 6), which shows the distributions for estimated β_2 s from four different methods: OLS, panel data FE, panel data RE and RCM estimation.

As with Figure 5, Figure 6 shows a wider distribution of estimates for OLS than the one obtained through other methods. Estimates obtained through FE and RE behave very similarly but their distribution is wider than for those obtained using RCMs, which translates into instances of significance (e.g. rejection of the null) even when the true parameter is zero. The narrower distribution with RCMs implies that more estimates of β_2 from among the 100 random samples are close to zero and therefore lacking statistical significance.

Importantly, random coefficient models in columns (8) and (10) also estimated very accurately the standard deviation of β_2 with an estimated coefficient of 0.492 in both cases, very close to the actual value of 0.5. A Strategy scholar facing these results would realize that, even when the average effect is not statistically significant, the variable is still strategically significant: for some firms the effect is positive; for others it is negative. This differential, firm-specific effect in performance may result in differences in competitive advantages that require further research.

Figure 7 exploits graphically this firm-level variation of the effect of X_2 on Y by showing the actual firm-specific coefficients for $\beta_{2,i}$ (in blue), their predicted coefficients $\hat{\beta}_{2,i}$ (in red), and

their 95 percent confidence intervals for first 50 firms in a random sample (sample 3).¹⁴ (We plotted only 50 observations to make the graph clearer.) Although the accuracy of the predicted values and their confidence intervals depends on the number of observations per firm, simulations available from the authors suggest that accurate estimates and tight intervals can be obtained with as little as five observations per firm for linear models. As suggested by the non-zero estimates for variance in models (8) and (10), the predicted firm-specific coefficients in Figure 6 show a diversity of firm-specific effects of X_2 on Y : for some firms it is positive, for some it is negative. For example, had X_2 been R&D investment, the researcher would have concluded that although the average effect of a dollar spent is zero in the sample, it is positive for some firms.

Finally, analyzing the results for β_3 provides insights on how RCMs can help Strategy scholars when the variation of an effect across firms is large. Recall that β_3 was generated from a random normal distribution with mean -5 and variance 25, which means that for some firms i in some samples, the parameter $\beta_{3,i}$ may be positive. Because the distributional mean for β_3 is far from zero, almost all models are able to reject the null hypothesis across samples. Even OLS is able to find statistical significance at the 5 percent level in 97 percent of the samples simulated. The estimates are also better across models; the farthest from the true value is OLS, but even there the differences are less striking than for the other independent variables. It is worth noting that, although results are similar for the models using FE, RE, and random coefficients, the latter are more accurate.

Figure 8 shows the distribution of the estimates for β_3 . As with β_1 and β_2 , an RCM produces the narrowest distribution around the mean of the true coefficient, -5.

¹⁴ Recall that these values are predicted for a given sample after the model has been estimated using Bayes theorem. This approach is therefore conditional on a specific set of parameters associated to a specific sample.

Besides precision, the value of using RCMs instead of traditional methods is evident by looking at the estimate of the standard deviation of β_3 . Under traditional methods, the Strategy scholar will assert that β_3 has a negative effect on Y_{it} . However, the large and significant variance of 5 suggests the possibility that, for some firms, the effect may in fact be positive. To illustrate, Figure 9 shows the predicted effect of X_3 for 50 firms within a random sample (again sample 3). Although most firms show a strongly negative effect, 23 show a positive effect, and one firm, firm 11, shows a strongly positive effect. Studying firm 11 more carefully might allow the researcher to develop new theories and hypotheses that advance our understanding of the phenomenon.

Finally, model (10) illustrates the estimation of covariance terms. Researchers using RCMs are able to estimate covariances among effects of independent variables. We simulated a negative correlation between the coefficients of β_2 and β_3 indicating that, for firms in which the marginal effect of variable X_2 is larger, the marginal effect of variable X_3 is smaller (and vice-versa). This suggests an implicit trade-off between X_2 and X_3 that is grounded on unobservable firm heterogeneity that simultaneously influences both effects—a correlation the RCM estimated very precisely.

RCMs estimated on Compustat data

Although the main features of RCMs are illustrated with our synthetic data analyses, we also examined its application using actual data. In choosing a setting to apply RCMs, we looked for several key features. First, the setting needed data that was widely available to academic researchers. Second, the data needed to be rich enough to reveal the key advantages of RCM estimation. Third, the questions that could be asked and answered with the data needed to be at the firm level and of interest to Strategy scholars.

After considering multiple sources, we decided to use financial data from Compustat to estimate a simplified log-linear Cobb-Douglas production function that explains firms' sales as a function of three production factors: capital, labor, and R&D. Specifically, we estimated the following equation:

$$SALE_{it} = \beta_i PPE_{it} + \gamma_i EMP_{it} + \delta_i XRD_{it} + \varepsilon_{it} \quad (8)$$

where S_{it} is firm-specific sales for firm i for year t ; PPE_{it} captures the value of net plant, property, and equipment for firm i in the year t ; EMP_{it} is the number of employees that firm i has in year t ; XRD is the amount invested in R&D by firm i during the year t ; and ε_{it} is a normally distributed iid error term with mean 0.

Compustat is ubiquitous in universities and most academics are familiar with it. It also contains thousands of firms across hundreds of industries, virtually assuring the presence of sufficient firm heterogeneity. Finally, we believe that models similar to those in equation (8) will be of interest for Strategy scholars. In fact, Knott (2012) uses a similar specification to calculate a firm-specific index that highlights firm-specific capabilities in R&D.

We downloaded from the Compustat Annual North America file data from 2001–2010, which provides a long enough panel to more accurately generate firm-specific estimates. The original dataset provided an initial 31,800 observations for 4,742 firms. After dropping any firm that had missing data for SALES, PPENT, EMP, or XRD—firm-year observations that would be lost in model estimation anyway—we were left with 19,563 observations for 2,932 firms. Finally, we dropped industries with less than 15 firms (734 firms, 4,997 observations) to assure a minimum level of heterogeneity within an industry, leaving a final sample with 14,586 observations for 2,198 firms in 43 industries.

We recognize that both our empirical model and measures could be improved. For the model, the dependent variable could be value-added instead of sales, we could add more variables (such as advertising expenditures) that may impact sales levels, etc. For the measures, some of these off-the-shelf data could be further manipulated to build variables that more accurately capture inputs into a production function. For example, annual R&D expenses should be used to construct a stock measure of R&D using a “perpetual inventory method,” where an initial stock is built and then, on an annual basis, stock is depleted at a set rate while yearly expense is added as new contributions. We instead chose simplicity, hoping that, by maximizing the reader’s ability to replicate our results, more would be encouraged to use this replication exercise as a springboard for their own applications using more fully articulated models.

With the data in hand, we searched for industries that illustrated the most important features of RCMs. We looked at a total of 43 industries before selecting two samples: Motor vehicles and passenger body (SIC=3711) and pharmaceutical preparations (SIC=2834). Table 3 shows the summary statistics per industry.

Table 4 shows the results for both industries for OLS, panel data fixed effects, random effects, and RCM estimation. Columns 1 to 4 correspond to motor vehicles, 5 to 9 to pharmaceutical preparations.

Results for PPE in the automobile industry illustrate the case of a variable that has a statistically significant effect on sales but it is not strategically significant because the variance of the effect is zero. In other words, all firms in the sample derive the same positive effect from a dollar spent on their plant and equipment, and variations in sales are simply a direct result of increases in capital investment levels. Assuming that access to capital is not a firm-specific

constraint, these results suggest that a firm may not be able to build a sustainable competitive advantage in this industry relying solely on plant and equipment investment.

The case of R&D investment in pharmaceuticals shows the ability of RCMs to discover that an effect that looks statistically insignificant is actually strategically significant, that is, of strategic interest. The coefficient for R&D offers mixed results: it is significant in FE but not in RE, suggesting the potential for substantial firm-specific heterogeneity in the effect. The RCM confirms that hypothesis. Although the mean effect under the RCM is not statistically different from zero, its variance is non-zero. This is not surprising given previous findings in the literature. For example, Henderson and Cockburn (1996) found large variation in R&D productivity among pharmaceuticals firms. In their case, they recognize this variation due to the large magnitudes and significance levels of firm dummy variables in a fixed effect panel model.

To further explore this firm-specific variation, researchers can predict firm-specific effects for each firm with RCMs and explore more deeply any variation observed. For example, Figure 10 shows the predicted firm-specific effects for the 10 largest firms (measured in terms of sales) in our sample. The figure shows that most firm-specific coefficients hover around 0 (e.g. the return on sales per dollar spent on R&D is practically negligible). However, some firms have a significantly positive effect (for example Merck, Astra Zeneca, and Aventis) while others show a negative effect (Pfizer, GlaxoSmithKline, and Roche).

We also introduced a model for pharmaceuticals that estimates correlations between the effects of all independent variables in column 9. Although the coefficient for the correlation between PPE and EMP is negative (suggesting a trade-off between capital and labor in the industry) the ones between XRD with any other explanatory variable are insignificant. A potential explanation for these patterns, which needs to be corroborated with more analysis, is

that these correlations likely document the differences between the business models of small biotech firms and big pharmaceutical companies. The former are known as the hotbed of innovation in the pharmaceutical industry. In contrast, large pharmaceutical companies are typically less innovative but have all the complementary marketing and distribution resources needed to successfully take new drugs to market. Thus, for smaller, more innovative biotech firms, R&D investments are expected to be marginally more effective at generating additional revenues through licensing, and employees should also be generally more productive at creating new, valuable intellectual property (IP) than in large pharmaceutical companies. This explains the positive correlation between the coefficients of R&D investment and number of firm employees. Also, since large pharmaceutical companies focus on the downstream part of the industry value chain—the production, distribution, and marketing of new drugs—their revenues are mostly driven by fixed capital investments and related efficiencies. To the extent that the opposite is true for small biotech firms, it explains the negative correlations between (a) the coefficients of PPE and R&D and (b) the coefficients of PPE and EMP. In other words, large pharmaceutical companies are more efficient in fixed capital investments, but R&D investments and employees are typically less productive. These correlations are essentially consistent with the revenue drivers of the two main business models in the pharmaceutical industry.

Summarizing, in this section we have showcased the advantages of RCMs using two examples: one with synthetic data generated to “fit” the conceptual underpinning of the model, and one with an actual, easy-to-access dataset. In both cases, the RCMs proved to be more precise in terms of coefficient estimation, while also providing researchers with a tool to explore and take advantage of the firm heterogeneity that is at the core of Strategy research. We believe that RCMs offer an exploratory and conclusive tool that can advance the development and

testing of theories in Strategy by unpacking heterogeneity, providing firm-specific effects, and allowing to model covariance in the effects of variables.

DISCUSSION AND CONCLUSIONS

In the last 30 years, Strategy scholars have made significant theoretical progress at understanding the differential effects of firms' actions on performance. Early foundations in Economics have been joined by diverse insights from other disciplines—sociology, political science, psychology, etc.—that have enriched our theoretic understanding of heterogeneity in firm performance. Yet, empirical work has not caught up with the growing sophistication of our theories. While in our models we theorize about firm actions that have differential effects on performance, empirically we only estimate the average effects of these actions across firms. Paradoxically, although our conceptual models are based on firm heterogeneity, we too often bury it in dummy variables and crude methods to obtain firm-specific marginal effects.

We believe that random coefficient models (RCMs) will help Strategy scholars to reduce the gap between theoretical and empirical research. RCMs assume that coefficients vary by individual, providing distributional estimates instead of fixed-point estimates. Specifically, RCMs deliver researchers two moments of the coefficients distribution: its mean (e.g. the average effect) and its variance (e.g. a shift parameter that allows us to pin down individual-specific effects).

RCMs provide three features that are particularly beneficial to Strategy scholars. First, RCMs allow us to model firm actions that have differential effects on performance. With variance estimates at hand, researchers can now make an important distinction between an explanatory variable that is statistically significant (its average effect is different from zero) and one that is strategically significant (its variance is different from zero). The latter term refers to variables

that, independently of whether (or not) they are statistically significant, have a firm-heterogeneous effect on performance, that is, where the effect of the variable on performance varies across firms.

We expect the spread of RCMs to foster new research in two directions. First, they can help reconcile prior, contradictory empirical findings that were exclusively grounded on average effects. For example, Anand and Byzalov (2012) demonstrated that more than 10 years of conflicting findings about the diversification discount (or premium) were rooted in firm heterogeneity: Although there is an average effect that supports the idea of a diversification discount, the effect varies by firm and becomes positive (indicating a diversification premium) for some firms. We believe there are multiple areas ripe for similar breakthroughs, for example, in the research on first-mover advantages. Second, we expect a wider use of RCMs to revive interest in explanatory variables that were previously written off as not worthy of further study, because their average effect across firms was found to be insignificant.

By producing firm-specific effects, RCMs offers researchers the ability to identify which firms are behind a differentiated effect. The ability to generate firm-level estimates means that any research context has the potential to become a teaching context. We can identify unique cases via firm-specific estimates and then follow up on these cases in more detail, which then might be used profitably in class. Teaching about unique cases in our own research settings might then cycle back and provide new research ideas. We can also easily imagine a different standard of research article that pairs large scale empirics with compact case studies of outlier firms identified through firm-specific estimates.

The final relevant feature of RCMs for Strategy scholars is the ability to model covariances between coefficients, providing a novel way to explore trade-offs between the actions firms take.

This approach complements existing tools such as the Pearson correlation of explanatory variables and interaction effects.

We demonstrated these three features using synthetic and actual data in examples that are easy to replicate and that can be used as training grounds for researchers interested in the RCM.

Even with all of our enthusiasm, a few caveats are in order. First, statisticians are still developing tools to perform sound statistical inferences for estimates of variance and covariance, so Strategy scholars should be conservative in claims of significance for these parameters. Second, an RCM is computationally demanding and its convergence is not guaranteed. This is likely to be the case when variances are close to zero or the models are complex (e.g. with many random coefficients and interaction among the explanatory variables).¹⁵ Finally, although this paper is heavily statistical in nature, we are Strategy researchers, not econometricians. In our own discovery, learning, and use of RCMs, we have focused on RCMs' application to our field. We have sketched the opportunities we see, but we are not experts in RCMs' assumptions and mechanics and cannot comprehensively delineate their uses and bounds. We expect greater clarity to emerge as researchers begin applying and refining the use of RCMs in Strategy. Our purpose here is to start that process.

We believe that Strategy is at a crossroads. On the one hand, we aspire to become a respectable field of study at par with the core disciplines we often benchmark ourselves against. On the other hand, a large fraction of our current research methods are unsuited to answer the questions that set Strategy apart as a distinct and unique field of research. This disconnect has never been formally addressed. We raise it here because we believe using RCMs appropriately in

¹⁵ The occurrence and solutions to these issues would vary depending on the statistical software. For example, see the Stata manual for tips on avoiding convergence problems.

future Strategy research would move the field a significant distance toward resolving this problem.

BIBLIOGRAPHY

- Anand B, Byzalov D. 2012. Systematic Heterogeneity versus Average Effects in the Returns to Diversification. Working paper, Harvard Business School.
- Besanko D, Dranove D, Shanley M, Schaefer S. 2007. *The Economics of Strategy*. John Wiley & Sons: New York.
- Chung W, Alcacer J. 2002. Knowledge Seeking and Location Choice of Foreign Direct Investment in the United States. *Management Science* **48**(12): 1534–1554.
- Cool K, Dierickx I. 1993. Rivalry, strategic groups and firm profitability. *Strategic Management Journal* **14** (1): 47–60.
- Craig S, Greene W, Douglas S .2003. Culture Matters: A Hierarchical Linear Random Parameters Model for Predicting Success of US Films in Foreign Markets. Working paper, New York University.
- Franco A, Sarkar MB, Agarwal R, Echambadi R, 2009. Swift and Smart: The Moderating Effects of Technological Capabilities on the Market Pioneering–Firm Survival Relationship *Management Science*. **55**(11): 1842–186.
- Ghemawat P. 2002 Competition and Business Strategy in Historical Perspective. *Business History Review* (76): 37–74
- Greene W. Convenient Estimators for the Panel Probit Model: Further Results. *Empirical Economics* (29)1: 21–48.
- Henderson R, Cockburn I. 1994. Measuring Competence: Exploring Firm Effects in Pharmaceutical Research. *Strategic Management Journal* **15**: 63–84
- Henderson R, Cockburn I. Scale, Scope, and Spillovers: Determinants of Research Productivity in the Pharmaceutical Industry. *RAND Journal of Economics* **27**(1): 32–59.
- Hoetker G. 2007. The use of logit and probit models in strategic management research: critical issues. *Strategic Management Journal* **28**(4): 331-343.
- Knott AM. 2008. R&D Returns Causality: Absorptive Capacity or Organizational IQ. *Management Science* **54**(12): 2054–2067
- Knott AM. 2012. R&D: The Trillion Dollar Fix. *Harvard Business Review*.
- Levinsohn J, Petrin A. 2003. Estimating Production Functions Using Inputs to Control for Unobservables. *Review of Economic Studies* **70**(2):317–342.
- Mas-Ruiz F, Ruiz-Moreno F. 2011. Rivalry within strategic groups and consequences for performance: the firm-size effects. *Strategic Management Journal*. **(32)** 12:1286–1308
- Nevo A. 2000. A Practitioner’s Guide to Estimation of Random-Coefficients Logit Models of Demand. *Journal of Economics & Management Strategy* 9(4): 513–548.

- Pacheco de Almeida, G, Hawk A, Yeung B. 2013. The Right Speed and Its Value. Working Paper Series, HEC Paris.
- Porter ME. 1980. *Competitive strategy*. Free Press: New York.
- Stata Multilevel Mixed-effects reference manual Release 13. Stata Press Publication, StataCorp LP, Colleges Station, Texas, 2013
- Topkis MD. 1998. Submodularity and Complementarity. Princeton University Press: Princeton, NJ.
- Train KE. 1998. Recreation Demand Models with Taste Variation over People. *Land Economics* **74**(2): 230–239.
- Revelt D, Train K. 1998. Mixed Logit with Repeated Choices: Households' Choices of Appliance Efficiency Level. *Review of Economics and Statistics* **80**(4): 647–657.
- Zelner B. 2009 Using Simulation to Interpret Results From Logit, Probit, and Other Nonlinear Models. *Strategic Management Journal* **30**(12): 1335–1448.

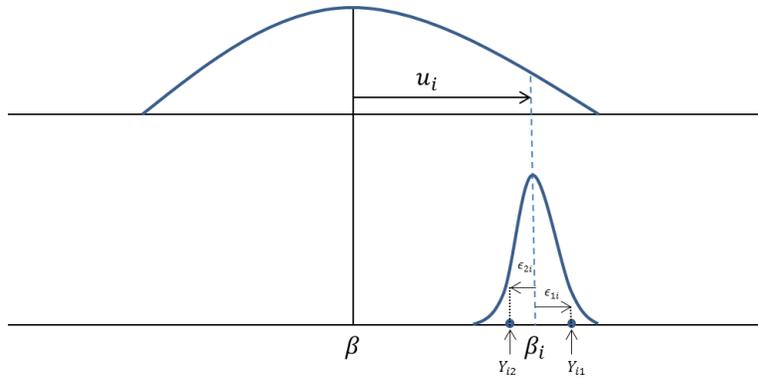


Figure 1¹⁶
Graphic representation for RCM

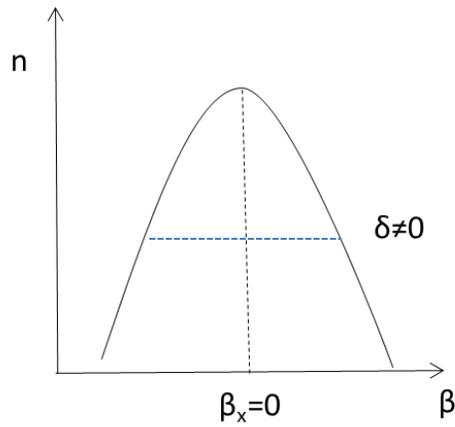


Figure 2
Case 1: statistically insignificant mean effect, variance different from zero

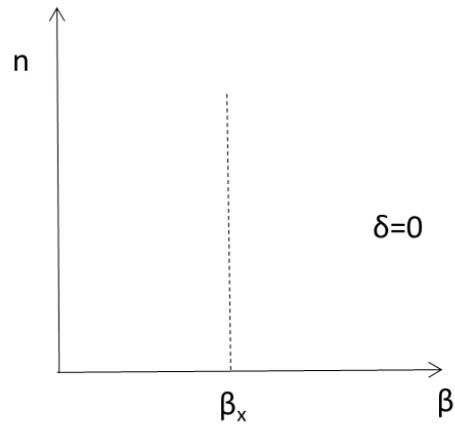


Figure 3
Case 2: statistically significant mean effect, zero variance

¹⁶ Figure from “Multilevel and Longitudinal Modeling Using Stata” by Rabe-Hesheth and Skrondal. 2005. StataCorp LP

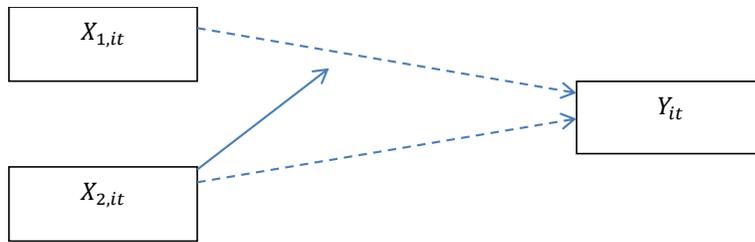


Figure 4
Interaction effect in Strategy

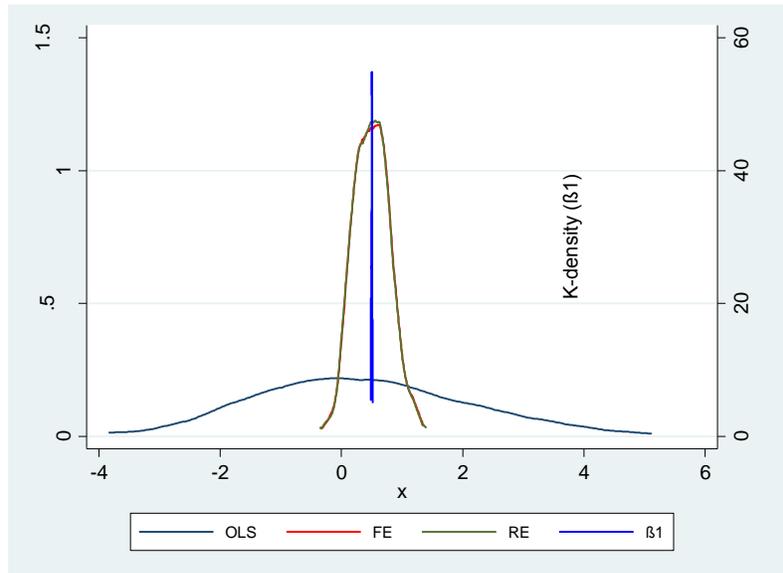


Figure 5
K-density of estimated parameters for β_1 – Multiple methods

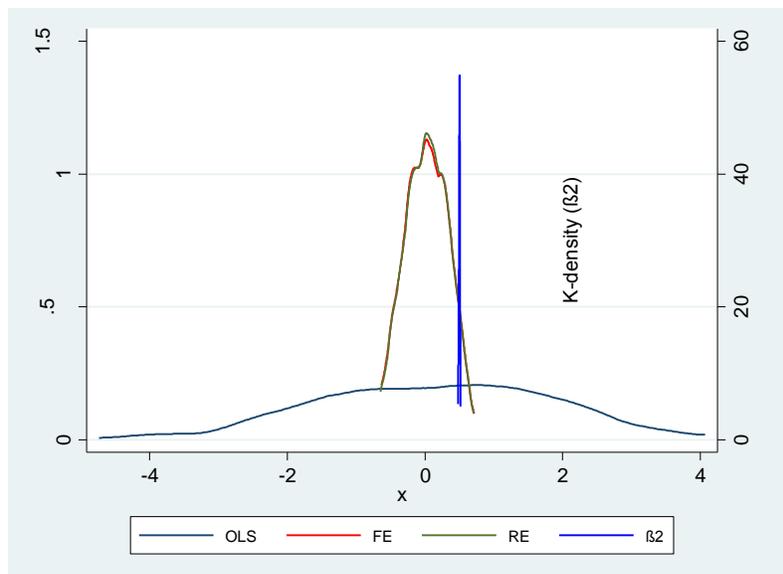


Figure 6
K-density of estimated parameters for β_2 – Multiple methods

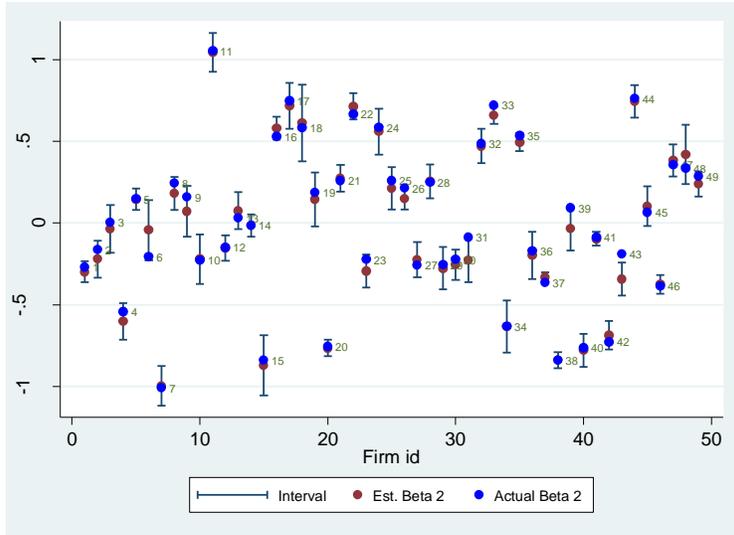


Figure 7
Predicted firm-specific estimates for β_2 – RCM

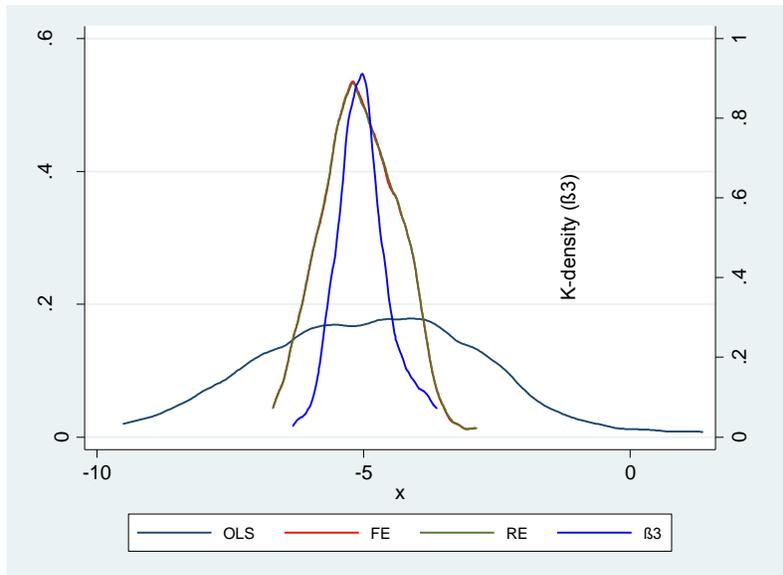


Figure 8
K-density of estimated parameters for β_3 – Multiple methods

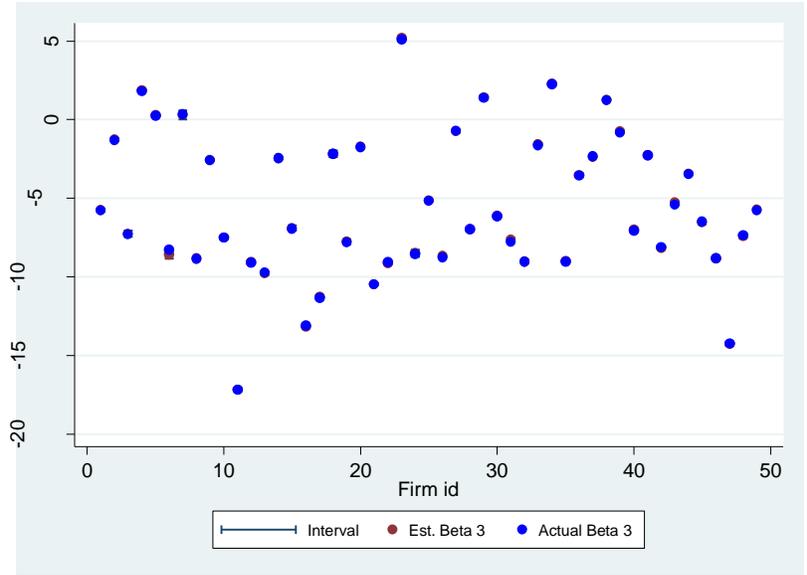


Figure 9
Predicted firm-specific estimates for β_3 – RCM

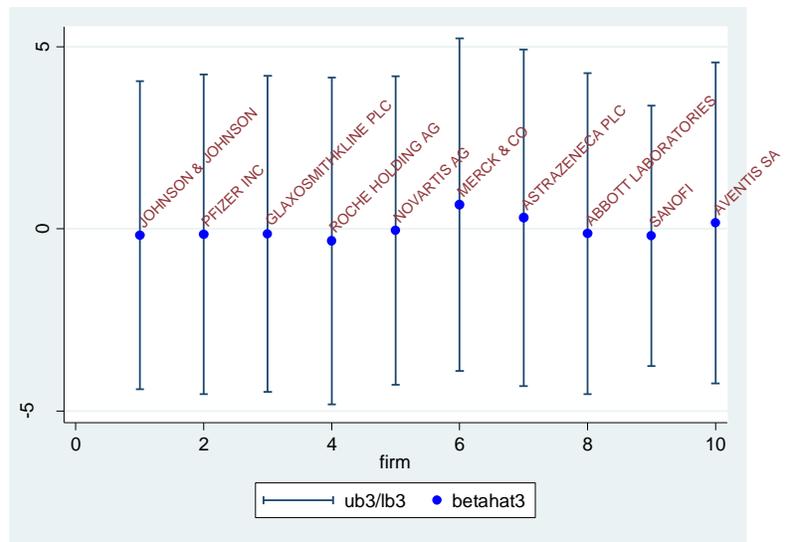


Figure 10
Predicted firm-specific estimates for XRD – Pharmaceutical industry

Table 1.a
Summary Statistics – Synthetic Samples

Variable	Obs	Mean	Std. Dev.	Min	Max
y	100,000	-223.80	255.19	-1712.68	918.30
x1	100,000	50.06	13.13	-14.85	109.90
x2	100,000	49.87	12.92	-10.05	114.16
x3	100,000	49.92	13.12	-10.28	113.95
beta1	100,000	0.50	0	0.50	0.50
beta2	100,000	0.00	0.50	-2.02	1.71
beta3	100,000	-5.00	5.00	-22.82	14.78
noise	100,000	0.00	1.00	-4.19	4.44

Table 1.b
Correlation Matrix – Synthetic Samples

	y	x1	x2	x3	beta1	beta2	beta3	noise
y	1							
x1	0.018	1.000						
x2	0.005	0.003	1					
x3	-0.250	-0.010	0.002	1				
beta1	1			
beta2	-0.386	-0.004	-0.004	-0.020	.	1		
beta3	0.926	-0.009	0.003	0.015	.	-0.504	1	
noise	0.004	0.003	-0.001	-0.004	.	-0.001	-0.001	1

Table 2
Summary of results for synthetic samples – Multiple methods

	True parameter	Cross-sectional		Panel data							
		OLS		FE		RE		RCM-var		RCM-cov	
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Coefficients											
β_1 (coefficient for X_1)	0.5	0.330	49%	0.501	70%	0.500	71%	0.501	100%	0.501	100%
		(0.594)		(0.165)		(0.164)		(0.008)		(0.007)	
β_2 (coefficient for X_2)	0	0.139	58%	0.016	34%	0.017	34%	0.002	7%	0.002	7%
		(0.604)		(0.165)		(0.164)		(0.050)		(0.050)	
β_3 (coefficient for X_3)	-5	-4.854	97%	-5.054	100%	-5.053	100%	-5.005	100%	-5.005	100%
		(0.594)		(0.164)		(0.164)		(0.496)		(0.496)	
Constant	0	-5.345	49%	2.715	47%	2.609	11%	-0.046	7%	-0.048	7%
		(52.383)		(14.226)		(28.167)		(0.635)		(0.631)	
Variance of coefficients											
St. dev. (β_1)	0							0.008	19%		
								(0.008)			
St. dev. (β_2)	0.5							0.492	100%	0.492	100%
								(0.035)		(0.035)	
St. dev. (β_3)	5							4.958	100%	4.962	100%
								(0.351)		(0.351)	
Corr (β_2, β_3)	-0.5									-0.503	100%
										(0.075)	
St. dev. (Residual)								0.995		0.998	
								(0.025)		(0.025)	
Observations		1000		1000		1000		1000		1000	
Number of firmid		100		100		100		100		100	
R-squared		0.09		0.52							
Log Likelihood		-6,902		-4,636		-5,026		-2,309		-2,294	
Improvement in Model Fit Test											
comparing				vs (1)		vs (1)		vs (1)		vs (1)	
difference in Log Likelihood				2266.70		1876.69		4593.54		4608.28	
additional d.o.f.				99		1		3		3	
Chi-square test of DLogL				0.000		0.000		0.000		0.000	

Table 3
Summary Statistics – Pharmaceutical and automobile industries

Variable	Obs	Mean	Std. Dev.	Min	Max
ln(sales)	1,890	3.73	3.17	-6.91	11.12
ln(PPE)	1,890	2.34	3.14	-6.91	10.03
ln(EMP)	1,890	-1.34	2.37	-6.91	4.81
ln(XRD)	1,890	2.91	2.57	-6.91	9.41
# of firms	304				
Automobile industry					
Variable	Obs	Mean	Std. Dev.	Min	Max
ln(sales)	195	8.93	3.29	-1.61	12.48
ln(PPE)	195	7.50	3.51	-2.98	11.45
ln(EMP)	195	3.08	2.77	-4.61	5.99
ln(XRD)	195	5.54	3.25	-6.91	9.30
# of firms	25				

Table 4
Results for Compustat data – Multiple methods

	Motor Vehicles and Passenger Car Bodies - SIC=3711				Pharmaceutical preparations - SIC=2834				
	OLS	FE	RE	RCM-var	OLS	FE	RE	RCM-var	RCM-cov
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PPE	0.0126 (0.062)	0.2138*** (0.058)	0.1629*** (0.058)	0.0749* (0.040)	0.037 (0.030)	0.1457*** (0.039)	0.1233*** (0.035)	0.1683*** (0.047)	0.1099* (0.058)
EMP	1.0844*** (0.073)	0.4699*** (0.114)	0.8175*** (0.091)	0.5271*** (0.136)	1.1817*** (0.040)	0.6842*** (0.067)	0.9402*** (0.053)	0.7872*** (0.062)	0.9102*** (0.085)
R&D	0.0669* (0.038)	0.1111*** (0.031)	0.1087*** (0.031)	0.1851 (0.222)	-0.0385** (0.019)	0.0996*** (0.037)	0.0369 (0.028)	0.0418 (0.033)	0.0394 (0.035)
Constant	5.1256*** (0.217)	5.2642*** (0.323)	4.4832*** (0.264)	4.2321*** (0.388)	5.3396*** (0.136)	4.0140*** (0.200)	4.4751*** (0.181)	3.9316*** (0.195)	4.1814*** (0.273)
stdv(PPE)				5.940E-14 (0.000)				0.390*** (0.037)	0.595*** (0.056)
stdv(EMP)				0.4877*** (0.102)				0.314*** (0.037)	0.895*** (0.084)
stdv(XRD)				1.0581*** (0.222)				0.223*** (0.041)	0.251*** (0.064)
stdv(constant)				1.2769*** (0.304)				0.867* (0.123)	2.8176*** (0.285)
corr(PPE,EMP)									-0.713*** (0.067)
corr(PPE,XRD)									-0.148 (0.212)
corr(PPE,constant)									-0.695*** (0.084)
corr(EMP,XRD)									0.102 (0.205)
corr(EMP,constant)									0.901*** (0.028)
corr(XRD,constant)									0.037 (0.214)
sd(Residual)				0.1495067 (0.010)				0.7981*** (0.018)	0.742*** (0.018)
Observations	195	195	195	195	1890	1890	1890	1890	1890
R-squared	0.973	0.967			0.803	0.205			
Log Likelihood	-158.3	2.927	-63.88	-50.45	-3326	-2387	-2917	-2839.1673	-2797.3736
F	2252	46.51			2561	135.8			
Number of firms		25	25	25		304	304	304	304