8-2013

Going Extreme: Systematically Selecting Extreme Cases for Study through Qualitative Methods

John A. DONALDSON
Singapore Management University, jdonaldson@smu.edu.sg

Ju Mae Rosie Ching
Singapore Management University, rosieching@smu.edu.sg

Andrew Tsao Ming TAN
Singapore Management University, andrew.tan.2009@business.smu.edu.sg

Going extreme:
Choosing exceptional cases for study through qualitative methods[1]

John A. Donaldson
Rosie Ching
Andrew Tan Tsao Ming
Singapore Management University

Delivered at the American Political Science Association Annual Meeting

Chicago, August 29-September 1, 2013

RESEARCH AND WRITING PRELIMINARY – PLEASE DO NOT CITE

Abstract

Oftentimes in social science research, we do not desire to study what is true on average. It may not be as helpful to know, for instance, that poor children from underperforming schools on average perform less well in college than wealthier children from well-managed schools. Especially if the variables in question are not readily amenable to change, it might be more useful to identify and closely examine exceptional cases. That is, studying why some poor children from underperforming schools do very well in college might identify some factors that could help raise performance levels despite unfavorable circumstances. However, the commonly used quantitative methods, including regression analysis, cannot help us do this.

Using a mixed-method approach, this paper focuses on how to apply quantitative methods to systematically select exceptional cases, which can be further examined through qualitative methods. To exemplify this method, the paper focuses on the relationship between economic growth and poverty reduction. While most economists agree that there is an inverse relationship between economic growth and poverty reduction, most also recognize that there are exceptions. From places that experienced rapid poverty reduction despite slow or negative economic growth, perhaps we can learn some lessons on what to do to make growth better for the poor. Similarly, from places that experienced slow or no poverty reduction despite rapid economic growth, perhaps we can learn some lessons on what *not* to do in our fight against poverty. Once we identify these cases, we can then turn to qualitative methods, including case studies and fieldwork, to trace the processes that caused these unexpected or puzzling results. While this does not deny the original relationship between economic growth and poverty reduction, it might be helpful for poor areas in which rapid economic growth is unlikely. This method can also be applied to other urgent social science questions.

[1] Parts of this paper are adapted from (Donaldson 2005; Donaldson 2008, 2011). The authors thanks Hal Wolman, Kris Ramsay, Kazuhiro Obayashi, Tomoko Fujii, Liu Tian, Luan Shenghua, Tom Holyoke and Aurobindo Ghosh.

Part I: Introduction

Much successful social science research has been conducted through discovering generalized relationships between variables. In political science for instance, the relationship between civil society and quality of democratic governance (Putnam 1993), the link between inequality and social conflict (Kanbur and Lustig 1999), and the relationship between economic growth and the emergence of democracy (Przeworski and Limongi 1997), have all been fruitfully investigated in part through regression analysis and other statistical methods designed to identify systemic relationships. By contrast, this paper bolsters the view that sometimes it can be just as important to identify and investigate exceptional cases – finding why some states with weak civil society have strong governance, why some highly inegalitarian societies display reasonable harmony, or why some wealthy states fail to become democratic. (Singapore could be an example of all three.)

Careful examination of such exceptions brings many benefits. It can often help us better understand the underlying dynamics that undergirded the original relationships and even identify previously undiscovered factors. Determining how such exceptions emerged can help suggest policy implications – positive exceptions can sometimes suggest ideas to adopt in other contexts, while negative exceptions can identify paths to avoid. Systematic selection of exceptions can also help avoid selection bias that tarnishes some qualitative studies.

This paper details and exemplifies a quantitative method for determining such cases. The authors are unaware of whether such a process has been detailed before – the use of this approach however does not seem to be common in the social sciences. It adopts a mixed method approach – the use of quantitative methods to identify exceptional cases than can be investigated through qualitative methods. Investigators must use this method carefully. It cannot be used in all situations. One must keep in mind that exceptions to relationships do not allow us to reject the original hypothesized relationship. On the other hand, some hypotheses seem to imply little hope for groups on the wrong side of the independent variable – implying that hope means beating the odds. Studying exceptions – those cases that beat the odds – can bring hope to such apparently hopeless situations. Perhaps the most useful applications of this method is that it can help us to identify new hypotheses to test and to help us identify classes of cases that defy exceptions.

Part II: Comparing approaches to case selection in qualitative research – a search for the exceptional

This method's success depends on selecting cases that are genuine exceptions to a given relationship. We must ensure that the cases are exceptions not due to random variation, but are instead caused by other more interesting (hopefully manipulable) factors not directly related to the hypothesis's independent variables. Previous studies on a range of issues have adopted many methods, some innovative, others expedient, to select cases. While such methods boast many advantages, they do not maximize confidence that the cases are indeed exceptions (see Table 1).

Method	Advantages	Disadvantages
Non-systematic methods (convenience, prior expertise)	Expedient	High risk of bias
Random selection	Minimizes bias related to human factors	Usually systematically selects typical cases
Consulting panel of experts (Cimadamore, Vidal et al 2002)	Increases number of involved experts from few to several	Risks following conventional wisdom; missing theoretically interesting but overlooked cases
Selecting cases with extreme values	Extremes likely to be exceptions; if phenomenon is important (though rare), can learn about causes	Change in dependent variable explained by commonly extreme values in commonly accepted explanatory variable; selection bias.
Typological design (Bennett and George 2002)	Finding exceptional cases	Presupposes extensive knowledge of cases.
Focused comparison (Lijphart 1971)	Controls for a number of factors	May not select true exceptions
Diverse cases (Gerring 2004)	Matches extreme values on both IV and DV	May not select true exceptions
Deviant cases (Gerring 2004)	Enriches understanding of relationship between IV/DV by focusing on true exceptions; enhances explanatory power over IV	Ability to generalize to other cases suspect

Table 1: A range of methods of selecting cases for qualitative studies

First, many social science studies focus on cases selected for practical reasons, including previous expertise in the area, or limited time or financial support. Examples abound. To some extent, most cases are selected because researchers have regional expertise, and utilizing it makes sense. While issues of time, money and expertise are significant, often unavoidable and even legitimate, cases chosen using non-rigorous methods are likely to be biased to some degree – leading to a common criticism of qualitative methods. A second method, random selection, chooses cases that are the opposite of what we seek; this method selects typical cases, not atypical ones. Moreover, random selection is simply not appropriate to subsequent qualitative research, because the risk of selecting a strange mix of cases increases as the number of cases decreases (Gerring 2004). Third, some scholars consult secondary research, or panels of experts to select cases, an innovative method that increases the numbers of experts focused on finding exceptions. For instance, Cimadamore and her colleagues (2002) empanel experts to help them select states within Argentina that exemplify 'best practices' for poverty reduction. However, this method risks following the conventional wisdom or adopting the biases of others, potentially missing theoretically interesting cases that experts have overlooked. A fourth option is to select cases with extreme values in their dependent variables, reflecting a high degree of success or failure. Selecting extreme cases of critical phenomena can be valuable on understanding the causes of rare, but important events. For instance, Theda Skocpol (1979) selected cases of revolution, an extreme and rare outcome, the study of which is inherently justified. However, if the extreme values are caused by equally extreme values on factors (potential independent variables) predicted by commonly accepted hypotheses, well established hypotheses may explain the phenomenon.[2] If so, choosing cases on extreme factors would not identify exceptions. A fifth alternative is to select cases through a typological design in which cases are divided into groups based on key criteria, and selected because they fall into theoretically interesting categories (Bennett and George 2002). Andrew Bennett et al (1994) use this methodology to select interesting cases based on the degree to which each contributed to the Persian Gulf War of 1991. This innovative method, although designed to select exceptional cases, often requires extensive knowledge on relevant factors for each case, which is often unavailable, or difficult or inefficient to obtain. Moreover, by combining cases into categories, this method is most applicable when data are in nominal or ordinal forms. If reliable integer or continuous data exist, however, typological design is probably not superior to the methods used in this study. Why divide a variable into categories – such

[2] Of course, this kind of study sometimes reveals that conventional wisdom is completely or partially wrong, in which case such a study is very useful. However, this is generally revealed after the case selection phase.

as high, medium and low – when they can be divided more accurately into continuous variables?[3]

The present study combines the three remaining case selection methods listed above. First, it follows the *most-similar-cases* design prescribed by Przeworski and Teune (1970), Lijphart (1971) and George (1979), selecting paired cases that are similar in as many exogenous factors as possible, thus controlling for them. By choosing cases that differ primarily in the factors of interest and are similar in many factors not of interest, this method increases rigor. Second, the study simultaneously focuses on *diverse cases*, in that though the background conditions for each case are similar, the variables of interest – the independent and dependent variables – diverge. Scrutinizing four cases (Korea, India, Brazil and Nigeria), Atul Kohli (2004) for instance selects across the range of the variables of interest, and includes those countries with different types of state structures and different types of results. In contrast to cases with extreme values, selecting diverse cases allows variance among factors in both the independent and dependent variables. However, neither selecting structured, focused comparisons cases nor selecting diverse cases ensures that the cases are indeed exceptional cases. Therefore, and most importantly, this process also focuses on *deviant cases*. Peter Evans, for example, focused on Brazil as a deviant case of a country that developed in spite of its dependence on international capital (Evans 1979). Deviant cases are particularly helpful for seeking "new – but as yet unspecified – explanations" (Gerring 2004, , p. 71), or at least credible hypotheses useful for testing in other contexts.

Part III: Describing the approach

[3] In a sense, the case selection method used here is a modified version of typological case design, using continuous data rather than nominal or ordinal data.

Choosing exceptional cases avoids many of the problems that tarnish other methods. After running regressions of data in ways that conform to the original hypothesis, we can select the points that are furthest away from the regression line, because those are the cases that are likely to be exceptions to the relationship. We can increase our confidence that the points furthest from the regression line are true exceptions by calculating the probability that any given point's position is random. For example, if a point is relatively near the regression line, the probability that its residual value (the distance away from the regression line) is generated by random factors is relatively high.

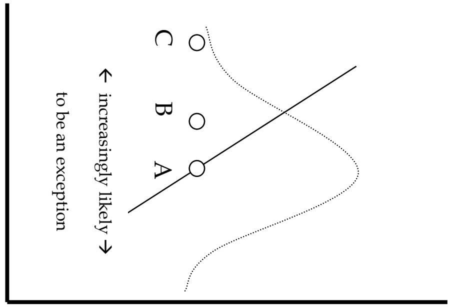

Figure: When can a point be considered an exception?

Exceptions to most regressions can be determined by computing the residuals, the vertical distance between the regression line and each data point, and then calculating the probability that the distance is random. Each point on the X-axis on the regression line corresponds to a bell-shaped curve that represents the chance that any given data point is random. This bell-curve has its modal value over the X-axis and extending in both directions along the Y-axis (see Figure above). For example, when a point (such as point "A") is positioned directly on the regression line, our confidence that such a point is explained by the independent variables is high. If the point is a modest distance (such as point "B") from the regression line, we can still be fairly confident that it is likely that the explanatory factors

explain that point, though our confidence decreases compared to point "A". As the distance between the point (such as point "C") and the regression line increases, our confidence that the explanatory factors explain that point diminishes. At some point, we will be willing to label such a distant point an 'exception' to the relationship. Calculating 'Z' scores can help us. By selecting cases with low 'Z' scores (say with a 'p' value below .05 in each direction), we can select the cases that are likely to be non-random exceptions. Choosing one case from each side of the regression line selects cases that defy expectations in both directions.[4]

However, even as it avoids many pitfalls, this method has at least three drawbacks. First, this method is sensitive to issues of data quality. A few misplaced data points will not normally affect regression analysis much, if the number of cases is large enough. If those flawed points are selected as exceptions, however, many months of qualitative research can be wasted in chasing a chimera. Since this step, case selection, is the first in a time-consuming and sometimes expensive larger research project, it is important to ensure that the data are reliable. Ideally, once the cases are determined, other descriptive statistics can be used to verify that they are correct. Second, it is possible that even true exceptions are created by systematic but theoretically uninteresting factors (uninteresting, that is, to social sciences), such as the environment (e.g., rainfall or earthquakes), geology (e.g., soil quality or insufficient water), or other non-policy omitted variable. The nature of the explanatory factors, moreover, would only be discovered after considerable research is completed. How disheartening it would be to find that the two exceptions were caused by some immutable or theoretically uninteresting factor. Similarly, sometimes incomparible cases can be selected through this process. Ideally, the case selection method described here will select cases that share several possibly significant, but theoretically uninteresting, factors, allowing us to control for them. We cannot always count on such a happy occurrence to provide a most-

[4] Using this method was a suggestion of Hal Wolman. The use of 'z' scores to test these probabilities was suggested by Kristopher Ramsay.

similar-cases design. This method can potentially pick cases that are so different from each other that they might be difficult to compare.

Third, selecting deviant cases can reduce potential application to other areas. Many contemporary political science studies generate generally applicable conclusions by discovering, with regression analysis, non-random associations between variables. To the extent that the researcher focuses on a random sample among a universe of cases, conclusions from such a study likely apply to many cases. However, this kind of study – even at its best – seeks general associations, and discards exceptions to discovered relationships between variables. Finding exceptions essentially does the opposite as this kind of analysis, using a regression not to seek the story in average cases, but to elicit exceptional cases, which may or may not represent a class to which multiple cases belong. This is based on a different type of approach – not that there is one general covering law, a goal to which regression analysis implicitly aspires (Goertz 2005). In the example below which focuses on the hypothesized relationship between economic growth and poverty reduction, the approach considers that there might be more than one pathway to poverty reduction. That in addition to economic growth, there are other ways to reduce poverty. Overall, the degree to which observations from these cases will be generalizable to others remains uncertain, and is largely a question for future research.

Although there are disadvantages of the process described here, selection bias is not one of them. Some might argue this does, because selecting cases with extreme values on the dependent variable can create a form of selection bias. As Collier and Mahoney argue, "selecting extreme cases on the dependent variable leads the analyst to focus on cases that, in predictable ways, produce biased estimates of causal effects," (Collier and Mahoney 1996, , p. 59). This case selection process is indeed guilty of one form of bias, because we are looking for atypical cases. The only problem with this, however, is that selecting on extreme values leaves us, as Collier and Mahoney note, less able to generalize our conclusions, a

point I discuss below. However, the case selection process, by selecting extreme cases from both ends of a range, does comply with a "basic and obvious rule: selection should allow for the possibility of at least some variation on the dependent variable" (King, Keohane, and Verba 1995, , p. 129). Selecting on the dependent variable, does not inherently bias the results by choosing cases that are known in advance confirm a favored hypothesis (for a discussion, see King, Keohane et al. 1995, pp. 128, 141-142), since at the time of the selection, the policies or detailed circumstances in the cases selected are not known.

Three advantages outweigh these three disadvantages. First, systematically selecting exceptional cases reduces the likelihood of choosing typical cases (as with random selection) or cases with certain forms of bias (through non-systematic selection methods). The use of quantitative data avoids the human element that other methods, such as empanelling experts, rely on, ones that could miss theoretically interesting unexpected cases. It avoids the problems of selecting cases with extreme values by controlling for theoretically uninteresting factors. Finally, this method's data requirements are often not as burdensome as those needed with the use of typological designs, which requires extensive qualitative knowledge of particular cases.

Second, by selecting cases through quantitative methods, and studying them through qualitative ones, we leverage the strengths of both tools. Most quantitative studies test theoretically interesting factors for statistically significant correlations. Instead, this approach results in a structured comparative case study, allowing us to collect a rich array of data directly from a variety of primary sources. In this way, even as the methods avoids some of the problems associated with case selection using qualitative methods, it also permits the deeper analysis of causes provided by such methods. Quantitative methods can be flawed. First, studies using overly aggregated statistics conceal provocative exceptions to even highly correlated relationships that qualitative methods can unpack (Ravallion 2001). Second, quantitative studies sometimes assume, rather than examine, causal linkages. Based on correlations, regressions often leave us unable to draw causal conclusions unless we supplement the method in some way, for instance, by examining underlying causal mechanisms.[5] Third, quantitative studies remove factors from their contexts, ignoring (or reducing to overly simplified variables) the history and background that are important to understanding these cases. Even if causality is established via quantitative methods, the mechanisms bringing about the relationship are concealed inside a black box – they generally have to be assumed or explored through other means. Qualitative research helps to illuminate the contents of that box. Rich information gleaned from carefully conducted fieldwork disaggregates averages, gets behind correlations to examine causal chains, permits the study of variables in their context and enhances our understanding of complex political phenomena, especially when cases are thoughtfully and rigorously selected (George and Bennett 2005). Therefore, this paper provides an example of how selecting case through quantitative methods enhances and supports qualitative research, without which an in-depth understanding of the dynamics of poverty change would not be possible.

Third, and most importantly, through identifying and scrutinizing exceptional cases, we can flesh out and modify the causal processes of old explanations, as well as potentially discover new, previously unidentified models. I exemplify this below.

Part IV: Example – Finding alternative pathways to poverty reduction through a study of two Chinese provinces

Scholars have previously identified economic growth as a major factor for reducing poverty in many cases and areas (Dollar and Kraay 2002). However, most economists also recognize that there are exceptions to this rule – many economies grow without poverty reduction and some see poverty reduction without economic growth (Donaldson 2008). Even if change in GDP explains about half of the changes in poverty – as per the results of

[5] For a discussion of causal mechanisms, see for example Levy and Goertz (2004).

one World Bank report (Dollar and Kraay 2002), we cannot be satisfied, given the importance and urgency of the problem of poverty. After all, many economies face barriers to growth. We need to find examples of such places that succeeded in reducing poverty despite those barriers. Moreover, there are multiple pathways to poverty reduction. By examining exceptions to the purported growth-poverty relationship, we might find each used one or a combination of these pathways, and thus learn more about the conditions under which these pathways did or did not work, and which specific mechanisms were used to reduce poverty (or not). Even more exciting, we might find that one of them traversed a new pathway that we can identify and detail. Identifying new pathways (such Peter Evans (1979), in the development context, did with Brazil) can illuminate new models that might help us understand poverty in the original case. To the extent that such models are examined and tested in other areas, our theoretical understanding of poverty also benefits. In terms of practical application, other polities can study from successful examples, and depending on original conditions, may be able to adapt judiciously some ideas in order to traverse a similar pathway.

One way to exemplify this is through my study of deviant cases by regressing changes in poverty rates against three factors that economists commonly believe cause poverty. Deviant cases are the points that are furthest away from the regression line, because those are the cases that are likely to be exceptions to the relationship between economic growth and poverty reduction. As obversered above, confidence that the points furthest from the regression line are true exceptions can be increased by calculating the probability that any given point's position is random. If a point is relatively near the regression line, the probability that its residual value (the distance away from the regression line) is generated by random factors is relatively high. That probability decreases when data points are further away from the regression line. Similarly, using 'z' scores we can calculate the probability of each point's residual value being random. By selecting cases with low 'z' scores (say with a ‘p’ value below .05 in each direction), we can select the cases that are likely to be non-random exceptions. Choosing one case from each side of the regression line selects cases that defy economists’ expectations in both directions: one in which poverty declined in a relatively stagnant economy, and another in which poverty persisted within the context of rapid economic growth.[6]

Selecting cases requires constructing a regression model using factors that reflect the expectations of economists as independent variables. Three factors are typically expected as being important: initial levels of poverty, geographic location and economic growth (see Table 2). First, economists expect that poverty in areas with higher poverty rates should change at rates different from those with lower rates. However, they disagree on that change’s expected direction, with some arguing that, other things being equal, poorer areas should become poorer over time (and richer areas richer), while others expect that convergence in poverty rates should occur through mechanisms such as spread effects.

	Variable	How calculated	Expected sign
DV	Change in poverty rates (DV)	Poverty rate: 1996 minus 1991	Not applicable (dependent variable)
IV1	Initial levels of poverty (IV1)	Poverty rates in 1991	Positive or negative
IV2	Growth rate per capita (IV2)	GDP/cap 1996-GDP/cap 1991, divided by GDP 1991	Negative (↑ growth → ↓ poverty)
IV3	Location (IV3)	Taken from Fan 1997	Coast (neg) > Inland > West (pos)

[6] Using this method was a suggestion of Professor Hal Wolman. The use of ‘z’ scores to test these probabilities was suggested by Professor Kristopher Ramsay.

Table 2: Variables used in regression

Second, economists expect that – except under relatively uncommon circumstances – economic growth reduces poverty rates: the faster the growth, the more rapid the decrease in poverty. Third, economists expect that, for historical reasons, regional differences in poverty rates should emerge. China scholars have detected a clear pattern of unequal development, with richer coastal areas outpacing the growth rate and development of central provinces, which in turn have outpaced the growth and development of western provinces. Moreover, some central government policies have been applied differently to the different regions – western, central and coastal (e.g., Fan 1997). Further, areas within these regions are similar, in the broadest terms, both geographically and demographically. For these reasons, generally, poverty in China’s coastal provinces should decline faster than poverty in inland provinces, which in turn should decline faster than poverty in western provinces.

With data collected from China’s National Bureau of Statistics (NBS) and the World Bank, I calculate a regression line based on a model (see Table 3) derived from three independent variables. As a dependent variable, I use changes in the poverty rates in 1991 and 1996 (disaggregated at the provincial level) calculated by World Bank economists based on NBS grouped data (World Bank 2001), the most reliable complete set of poverty data for each of China’s provinces that is currently available to the public. I excluded from the dataset three of China’s provincial-level cities (Beijing, Tianjin and Shanghai) because of their tiny rural populations. Since China’s fourth provincial-level municipality, Chongqing, did not become an independent provincial-level city until 1997, I am comfortable in including Chongqing’s data in Sichuan’s for 1996. In addition, fearing inaccurate data, I omitted two provincial-level autonomous regions, Tibet and Xinjiang. Because these two autonomous regions contain large independence movements with international support, China's central government has particular reasons for not reporting accurate growth and poverty data there. Eliminating these six provincial-level governments leaves a total "N" of 25 for each of the two years.

Source	SS	df	MS		
Model	2149.05211	4	537.263029	F(4, 20) =	20.99
Residual	511.869347	20	25.5934674	Prob > F =	0.0000
				R-squared =	0.8076
				Adj R-squared =	0.7692
Total	2660.92146	24	110.871728	Root MSE =	5.059
				Number of obs =	25

povt2t1intl	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
pov91intl	-.6909107	.0934758	-7.39	0.000	-.8858978	-.4959236
gdpcapgrwth	-.0506892	.0276951	-1.83	0.082	-.1084602	.0070818
coast	-10.47261	4.72723	-2.22	0.039	-20.33344	-.6117824
inland	-14.03452	3.205835	-4.38	0.000	-20.72177	-7.347266
_cons	25.87969	6.803142	3.80	0.001	11.68859	40.0708

Table 3: Results of regressing initial levels of poverty, growth and location against changes in poverty rates

If this were a hypothesis-testing exercise, this regression would be misspecified and inaccurate.[7] However, this exercise is not about testing the economic model, but is instead trying to be as fair to the model as possible. The model serves to predict, but not explain, as much of the variation in poverty rates as possible through non-policy factors. Since correlation between variables increases the likelihood of finding a significant result, using correlated factors as independent variables is both fair to the model and not inaccurate as long as we attach little importance to a significant finding. Moreover, it is difficult to imagine omitting one of the variables simply because another variable can serve as proxy,

[7] For instance, the independent variables are not truly independent of each other. GDP growth, we know from extensive research, has been fastest in China's coastal provinces, and indeed two variables (GDP growth and the costal dummy variable) have a .64 correlation. Similarly, initial levels of poverty should be related both to location (the poverty variable has a .62 correlation with the "west" dummy variable) and to GDP growth (with a .54 correlation).

since it would be less accurate a reflection of economists' expectations, even if it might come up with similar results.[8] Similarly, there is no need to test for heteroskedasticity since the model we are testing does not expect the results to be heterskedastic. Even if the variation varies across the entire range, it would have no bearing on finding residuals. Nor do I test for a non-linear relationship between the independent and dependent variables, since the model I am simulating argues that there is a one-to-one correlation between economic growth and poverty reduction. Most economists would expect a straight regression line.

The model as a whole is statistically significant (although, given the discussion above, little weight can be attached to this result). The first factor, initial levels of poverty, is significant at the .05 confidence level. The negative sign indicates that poverty declined more in provinces with higher initial poverty rates, consistent with the convergence hypothesis. The second factor, per capita GDP growth, is significant at a 0.1 confidence level, and has a negative sign, meaning that poverty rates in the provinces with higher rates of economic growth declined more than did poverty rates in provinces with lower rates of growth. This is consistent with the view that poverty in provinces with a higher growth rate per capita declines faster than it does in provinces with a low growth rate per capita. Location, the final factor, is similarly significant, with signs in the expected direction.

[8] For instance, location not only captures economic growth, but more importantly controls for central-level policy and spread-effects. Running the same regression without controls for location has similar results: statistically significant ($p<.0000$) with a smaller R^2 of .6. Although Yunnan remains the most extreme exception in one direction (in fact, the residual for Yunnan increases to 18) when location dummy variables are omitted, Guizhou is not chosen. Instead, the central province of Anhui, which experienced a high degree of poverty reduction, is chosen. However, Anhui borders to the two rapidly growing provinces, Jiangsu and Zhejiang, and the result can be superficially explained by spread-effects and migration (for which Anhui is well known), factors for which the location dummy variables at least partially control.

The model is also an accurate predictor of changes in poverty. The adjusted R^2 for the model is nearly 0.77, implying that the three non-policy determined factors selected explain more than three-quarters of the variance in the provincial poverty rates, although this is largely caused by the way the model was constructed. Despite the above discussion on model specification, stopping at this point is nevertheless tempting. Given such a close fit, the results of this model are consistent with the expectations of mainstream economists. Growth, combined with location and initial levels of poverty, helped to reduce poverty in China between 1991 and 1996.

A closer examination, however, reveals that while the changes in poverty rates of no less than six of the 25 provinces are within one percentage point of that predicted by the model – exceptionally accurate – and the majority (14) are within three points (see Table 4, below), the model has no less than four exceptions. These are identified through a process described below.

Province	Predicted poverty change	Reported poverty change	Residual (Reported - Predicted)	"Excep-tions"	Difference in human terms (Residual x Rural pop)
Guizhou	-21.94	-30.83	-8.89	0.03	-2,650,285
Sichuan	-11.01	-19.21	-8.20	0.04	-7,614,187
Anhui	-44.57	-50.84	-6.27	0.09	-3,116,533
Guangdong	-0.92	-5.58	-4.66	0.16	-2,233,282
Jiangsu	-11.37	-14.55	-3.18	0.25	-1,627,759
Hebei	-19.65	-22.74	-3.09	0.25	-1,642,810
Jiangxi	-9.04	-11.77	-2.73	0.28	-856,231
Liaoning	-5.23	-7.51	-2.28	0.31	-510,861
Shandong	-11.67	-13.02	-1.35	0.39	-875,742
Henan	-34.47	-35.17	-0.70	0.44	-534,195
Ningxia	-13.16	-13.66	-0.50	0.46	-18,929
Hubei	-20.02	-20.32	-0.30	0.47	-125,034
Zhejiang	-6.84	-7.02	-0.18	0.48	-66,009
Fujian	-11.32	-11.23	0.09	0.51	24,385
Hunan	-14.10	-13.77	0.33	0.53	173,757
Heilongjiang	-18.77	-17.69	1.08	0.59	214,620
Qinghai	-12.73	-11.4	1.33	0.61	44,264
Jilin	-14.24	-12.54	1.70	0.64	252,053
Shaanxi	-14.60	-12.45	2.15	0.68	586,826
Shanxi	-28.03	-25.58	2.45	0.70	562,961
Gansu	-27.22	-23.74	3.48	0.77	690,920
Inner Mongolia	-23.34	-18.91	4.43	0.83	672,480
Guangxi	-18.60	-11.69	6.91	0.93	2,606,050
Hainan	-12.13	-4.38	7.75	0.95	418,950
Yunnan	-15.83	-5.19	10.64	0.98	3,578,028

Table 4: Provinces, in order of residuals

I determine the exceptions to this regression by computing the residuals, or the distance of each point from the regression line. Calculating "z" scores to this estimates it is the probability or chance of getting such a datum if the standardized residuals obey the standard normal distribution. The smaller the *p*-value is, the more unlikely it is to get such an observation under the standard normal curve, thereby giving rise to the conclusion of it being an exceptional case in the direction of which tail wherein it lies. When the chance of that is less than five percent in either direction, that point is labeled as an exception. Specifically, I subtract the change in poverty rates predicted by the model (column 2 in Table 4) from the value reported by the data (column 3 in Table 4). This "residual" value is recorded in column 4 in Table 4. Column 5 (labeled 'exceptions') records the probability that the position of each province's corresponding data point is caused by random variation. This is estimated by calculating the "p" values of the residual's "z" scores to determine the chance that the point's distance from the regression line is caused by random factors (the "p" values are in column 5). As noted above, I consider any province with a "p" score higher or lower than five percent to be considered an exception. Thus, provinces with an "exception" score less than .05 and greater than .95 are considered exceptions (an equivalent "p" value of .1 in a two-tailed test). These four provinces I have highlighted in gray on Table 4. The final column (differences in human terms), calculated by multiplying the residual times the rural population of each province, expresses the difference between the predicted value and the actual value of rural poverty. In other words, this column is an attempt to estimate how many more or fewer people would be in poverty in each province if economic growth had had the effect on poverty that the model predicted (i.e., all data points were on the regression line).

The regression therefore suggests that four provinces are exceptions in China. Although the model does explain the overall pattern of poverty reduction in China, for the exceptions, the difference between the model's expectations and actual change is significant (as calculated by "z" scores) and considerable. Between 1991 and 1996, Guizhou and Sichuan's poverty reduction outperformed what the model predicts based on the three independent variables. Guizhou's poverty rate declined 8.89 more percentage points than the model predicts, meaning that an additional 2.65 million people were not poor in 1996 than would be true if the model accurately predicted that point. Similarly, Sichuan's poverty rate declined 8.2 percentage points more than expected, meaning that an addition 7.6 million people emerged from poverty.[9] Yunnan and Hainan were exceptions in the other direction: their poverty rates declined far less than the model predicts. Yunnan saw a 10.64 percentage point difference between the predicted and actual poverty rates, implying an additional 3.58 million people were in poverty in 1996 than the model predicts. Hainan had a difference of 7.75 percentage points, translating to about 419,000 people.

Part V: Concluding thoughts

From here, the research would use typical qualitative methods, including interviewing, surveys, fieldwork observation and analysis of descriptive statistics, to test possible reasons why the cases were exceptions to the proposed hypothesis. The methods of conducting these qualitative approaches have been extensively described elsewhere (e.g., George and Bennett 2005). The fact that there could be dozens or more possible explanations for the puzzle

[9] The World Bank reports Chongqing's data separately from Sichuan in 1996. Because Chongqing did not formally separate from Sichuan until March of 1997, I added Chongqing's GDP and poverty data to Sichuan's before running the regression.

identified demands considerable work. However, if this results in possible alternative pathways to poverty reduction, these endeavors are worth it.

Overall, this methodology is in the spirit of World Bank economist Martin Ravallion's (2001) argument "People are often hurting behind the averages. Panel data and observations from the ground can reveal this, but the aggregate statistics cannot. It is important to know the aggregate balance of gains and losses, but it will be of little consolation to those suffering to be told that poverty is falling on average,'' (Ravallion, 2001, p. 1811)." It also reflects economists Michael Lipton and Robert Eastwood's argument, "Faster growth is normally better for the poor than slower growth, and is not systematically offset by any change in distribution. But huge exceptions – and the possibility of clusters of countries where growth is much better for distribution or much worse – mean that these findings are the beginning, not the end, of the inquiry. Residuals matter," (Eastwood and Lipton 2001, p. 16). These sentiments apply not only to research into poverty, but to a range of other social science questions.

In typical regression analyses, residual cases are dismissed. In this study, they are of central importance. From a theoretical perspective, while these exceptions cannot allow us to reject the original hypothesis, studying them can help us to understand theorized relationships more deeply, and to see perhaps classes of exceptions. For real world practicality, residuals can matter in a number of issues of vital importance. From positive cases like Guizhou, perhaps we can find some suggestions on what to do when rapid growth is unlikely. Negative cases might illuminate some pitfalls – ways that robust economic growth does not translate directly to poverty reduction. When poverty in endemic, the search for exceptions is important. Where war is likely, the search for peace in spite of the odds can save lives. Where oppression and civil strive are the rule, finding alternative pathways can promote harmony.

Bibliography

Bennett, Andrew, and Alexander L. George. 2002. *Developing and Using Typological Theories in Case Study Research* 2002 [cited February 16 2002]. Available from http://www.georgetown.edu/bennett/PROTCG.htm.

Bennett, Robert J. 1994. *Local Government and Market Decentralization: Experiences in Industrialized, Developing, and Former Eastern Bloc Countries*. Tokyo and New York: United Nations University Press.

Cimadamore, Alberto, Erika Vidal, Fabiana M. Werthein, and Michel Flament. 2002. "Best Practices in Poverty Reduction in Argentina: Towards the Identification of a Selection Method." In *Best Practices in Poverty Reduction: An Analytical Framework*, edited by Else Øyen, 88-107. New York, NY: St. Martin's Press.

Collier, David, and James Mahoney. 1996. "Insights and Pitfalls: Selection Bias in Qualitative Research." *World Politics* no. 49 (1):56-91.

Dollar, David, and Aart Kraay. 2002. "Growth is Good for the Poor." *Journal of Economic Growth* no. 7 (3):195-225.

Donaldson, John. 2005. *The Political Economy of Local Poverty Reduction: Economic Growth, Poverty Reduction and the State in Two Chinese Provinces*. Ph.D. Dissertation, Department of Political Science, The George Washington University, Washington DC.

Donaldson, John A. 2008. "Growth is Good for Whom, When, How? Economic Growth and Poverty Reduction in Exceptional Cases." *World Development* no. 36 (11):2127-2143.

———. 2011. *Small Works: Poverty and Economic Development in Southwestern China*. Ithaca NY: Cornell University Press.

Eastwood, Robert, and Michael Lipton. 2001. Pro-Poor Growth and Pro-Growth Poverty Reduction: What Do they Mean? What Does the Evidence Mean? What Can Policymakers Do? Manila: Asian Development Bank.

Evans, Peter B. 1979. *Dependent Development: The Alliance of Multinational, State, and Local Capital in Brazil*. Princeton: Princeton University Press.

Fan, C. Cindy. 1997. "Uneven Development and Beyond: Regional Development Theory in Post-Mao China." *International Journal of Urban and Regional Research* no. 21 (4):620-639.

George, Alexander L. 1979. "Case Studies and Theory Development: The Method of Structured, Focused Comparison." In *Diplomacy*, edited by Paul Lauren, 43-68. New York: Free Press.

George, Alexander L., and Andrew Bennett. 2005. *Case Studies and Theory Development in the Social Sciences*. Cambridge: MIT Press.

Gerring, John. 2004. *Case Study Research: Principles and Practices (draft version: November 26, 2004)*.

Goertz, Gary. 2005. Social Science Concepts: A User's Guide.

Kanbur, Ravi, and N. Lustig. 1999. Why is Inequality Back on the Agenda? Paper read at Annual Bank Conference on Development Economics, at Washington, DC.

King, Gary, Robert O. Keohane, and Sidney Verba. 1995. *Designing Social Inquiry*. Princeton: Princeton University Press.

Kohli, Atul. 2004. *State-Directed Development: Political Power and Industrialization in the Global Periphery*. Cambridge: Cambridge University Press.

Lijphart, Arend. 1971. "Comparative Politics and the Comparative Method." *American Political Science Review* no. 65:682-693.

Przeworski, Adam, and Fernando Limongi. 1997. "Modernization: Theories and Facts." *World Politics* no. 49 (2):155-183.

Przeworski, Adam, and Henry Teune. 1970. *The Logic of Comparative Social Inquiry*. New York: Wiley-Interscience.

Putnam, Robert. 1993. *Making Democracy Work: Civic Traditional of Modern Italy*. Princeton: Princeton University Press.

Ravallion, Martin. 2001. "Growth, Inequality and Poverty: Looking Beyond the Averages." *World Development* no. 29 (11):1803-1815.

Skocpol, Theda. 1979. *State and Social Revolutions: A Comparative Analysis of France, Russia and China*. Cambridge: Cambridge University Press.