

NEPC Review: School Sector and Satisfaction: Evidence from a Nationally Representative Sample (Annenberg Institute for School Reform at Brown University, October 2019)



Reviewed by:
Steven V. Miller
Clemson University

January 2020

National Education Policy Center

School of Education
University of Colorado Boulder
nepc.colorado.edu

Acknowledgements

NEPC Staff

Kevin Welner
NEPC Director

William Mathis
Managing Director

Alex Molnar
Publications Director


Suggested Citation: Miller, S.V. (2020). *NEPC Review: “School Sector and Satisfaction: Evidence from a Nationally Representative Sample.”* Boulder, CO: National Education Policy Center. Retrieved [date] from <http://nepc.colorado.edu/thinktank/school-satisfaction>.

Funding: This review was made possible in part by funding from the Great Lakes Center for Educational Research and Practice.



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

This publication is provided free of cost to NEPC’s readers, who may make non-commercial use of it as long as NEPC and its author(s) are credited as the source. For inquiries about commercial use, please contact NEPC at nepc@colorado.edu.



NEPC Review: School Sector and Satisfaction: Evidence from a Nationally Representative Sample (Annenberg Institute for School Reform at Brown University, October 2019)


Reviewed by:

Steven V. Miller
Clemson University

January 2020

Executive Summary

School Sector and Satisfaction: Evidence from a Nationally Representative Sample, a recent report published by the Annenberg Institute for School Reform at Brown University, argues that private and charter schools have a strong effect on parents' reported satisfaction with their children's education. It explains this effect by pointing to competitive pressures and the importance of reducing the monopoly power of the traditional public school system. The report's analyses are based on a nationally representative sample from the National Household Education Surveys Program (NHES). However, the report suffers from two major flaws. First, it saturates the important analyses with over 230 covariates. This amounts to a "garbage can" regression modeling approach that obscures more than it illuminates; variables are misspecified and results are sensitive to the oversaturation of the regression model. A reader versed in statistical modeling will have no confidence in the substance of the findings. Second, and even more importantly, the report's decision to focus on just the "very satisfied" overstates the effect of private and charter schools on parent satisfaction. Almost 90% of public school parents are satisfied with their child's education, and the report's decision to focus on just the "very satisfied" appears to be a deliberate modeling choice to greatly overstate the purported effects of private and charter schools on parent satisfaction. This appearance is confirmed by the report's problematic use of past research literature, again suggesting an interest in findings that will support advocacy for curtailing the supply of public education. For each of these reasons, the report is of little or no use to policymakers and others with an interest in understanding parent satisfaction associated with school choice.



NEPC Review: School Sector and Satisfaction: Evidence from a Nationally Representative Sample (Annenberg Institute for School Reform at Brown University, October 2019)

Reviewed by:
Steven V. Miller
Clemson University

January 2020

I. Introduction

The report starts with an economic interpretation of the U.S. public school system, arguing that one of the foundations of American social life enjoys considerable monopoly power. A child's public school is residentially assigned, and not chosen by the child or the child's parents. This creates a conundrum for parents who are dissatisfied with their child's public school since the report contends that recourse is either costly (e.g., paying for a private school while still subsidizing the public school through taxes) or ineffective (e.g., advocating for the child's well-being in the public school). Public education is a public good, but one interpretation sees education providers with far more power than education consumers.

Increasing access to alternative schooling options offers one path toward increased satisfaction for parents with their child's education. Theoretically, parents with access to an alternative schooling option better resemble "customers" in the marketplace of education services, and they have more control over the service provided to them. These alternative schooling choices are, per this interpretation, more responsive than public schools to parent feedback because they operate on market terms. This should increase parent satisfaction with the educational product. The report's analyses of the *NHES Parent and Family Involvement in Education* survey, a nationally representative survey of 13,436 respondents, provides evidence that suggests parents are more satisfied with charter schools and private schools than public schools on a variety of dimensions, and that their child seems happier for the experience.

This review briefly summarizes the report's findings and the audience it addresses. Thereafter, the review explores two major limitations with the analyses provided in the report.

These include the “garbage can” regression modeling approach and the selective interpretation of how satisfied parents are with public schools vis-à-vis other schooling options. This review offers a reanalysis of the central claims made by the report, suggesting only modest benefits of charter schools and private schools on parent satisfaction in relation to public schools.

II. Findings and Conclusions of the Report

The 2016 NHES includes items on parent satisfaction with the child’s school, what type of school the child attends in addition to various sociodemographic indicators of interest (e.g., parent’s age, education, income, race/ethnicity and the size and nature of the child’s family).

Table 1 summarizes a series of *t*-tests that look for differences in means among a variety of dimensions among public school parents, private school parents, and public charter parents. Of note, public charter parents and private school parents were discernibly more likely to be very satisfied with their child’s school overall, the child’s teachers, the academic standards, the order and discipline of the school, and staff interaction compared to public school parents. Public charter and private school parents are also more likely than public school parents to strongly agree that their child seems to be enjoying the school..

Tables 2 through 7 serve as the core findings of the report. The results here are akin to a “fixed effects” analysis, in which the child’s school type—private (Catholic), private (other religious), private (non-religious), public charter, public—is broken into a series of dummy variables that isolates the effect of a particular school type on one of the six items of interest relative to a baseline comparison (here: traditional public schools). Each table includes eight models that gradually add more control variables, amounting to 236 variables in the final model for a given outcome variable. The coefficient for the public charter dummy variable and the private school variables are positive and statistically significant in every application, suggesting that charter school parents and private school parents report more satisfaction than public school parents.

The final four tables look for the effects of alternative school types on parent satisfaction by the parent’s education level (Table 8), the parent’s income level (Table 9), the student’s race/ethnicity (Table 10), and whether the student has a disability (Table 11). The results here suggest that public charters and private schools have robust effects on parent satisfaction at different levels of the parent’s income level, student race/ethnicity, and the presence or absence of a disability for the child. However, the effects of public charters do not appear as robust on the item that asks the parent whether the child seems to be enjoying the school.

III. The Report’s Rationale for Its Findings and Conclusions

The report’s rationale for the findings is both theoretical and empirical. Theoretically, charter schools and private schools provide an alternative schooling option for parents who may

be dissatisfied with the near monopoly power of the American public school system. These schools should be more responsive to parent demands, better empowering education consumers. A parent's reported satisfaction should therefore increase. Empirically, the data are nationally representative and sufficiently powered. At first glance, a statistically significant effect of the school type dummies compared to public schools, when almost every other indicator in the data is included as a statistical control, should suggest the positive effects of alternative schooling options on parent satisfaction are not sensitive to omitted variable bias.

IV. The Report's Use of Research Literature

A reader will come to two conclusions about the research literature in the report. First, the report is clear from its opening paragraph about the narrow view the author holds about public education; he advocates for defunding public schools to pursue alternative schooling options for a narrow population. The report eschews classic works in comparative public policy that see public education as a public good and a longstanding government tool for the promotion of the general welfare.¹ The report instead offers a constricted view of this public good as a monopoly. It then follows the logical implications of how monopolies constrict consumer agency, thus characterizing public education, which is a public good, as inherently bad. In just the first paragraph, the report communicates that it is uninterested in other perspectives on public goods and the historical purposes of public education in the United States. Empirical issues aside, the use of past research literature to start the report suggests the report is only interested in a certain set of findings that underscore advocacy for curtailing the supply of public education. The reader will know from the first sentence of the report that this is less about research and more about advocacy.

Second, the past research discussed more extensively in the literature review runs into similar problems. The literature reviewed tends to produce findings of value to those advocating school choice measures. Publication bias will be a nagging concern in the school choice literature, especially for the *Journal of School Choice* dedicated to research on this topic. The bias is toward studies that report positive findings in line with school choice advocacy. There is the familiar problem of overemphasizing the lottery studies while underemphasizing the quasi-experimental studies. The lottery studies lack important external validity because of self-selection into the lottery, but the results are more in line with the intended aims of this report and others like it. The quasi-experimental studies do report some positive effects of school choice but the findings are not as sanguine as the lottery-based studies the report does review. Indeed, a review of those quasi-experimental studies will offer far more caveats about the effectiveness of school choice than a reader would discern in this report.²

V. Review of the Report's Methods

A reader versed in regression modeling and the statistical analysis of survey data will be drawn to two major shortcomings in the report's methodology. The first shortcoming concerns how the report handled potential confounding effects; an oversaturated model will obscure more than it will illuminate. The second shortcoming concerns the selective inter-

pretation of the dependent variables of interest and the overemphasis the report places on the “very satisfied” and “strongly agree” responses.

A ‘Garbage Can’ of Confounding Effects

The most glaring shortcoming in the report concerns how it addressed potential confounding effects. “Confounding” effects, colloquially, are effects that could influence the results of interest (i.e., the effect of school type on parent satisfaction), but are not included in the model as statistical controls. Table 1 highlights that there are substantial differences among public school parents and charter and private school parents on dimensions like race/ethnicity, educational attainment of the parent, and characteristics like household income and public assistance. If these confounding sources are excluded from the statistical analysis, the effect of charter schools is potentially biased. Thus, it would be wise to isolate these potential sources of confounding as the report tries to do.

However, the report’s handling of statistical controls is not advisable. 236 independent variables in a regression model is conceptual overkill by any measure. Social scientists adept at this type of analysis would dismiss this approach as a “garbage can model.” Simply, the signal is overwhelmed by the noise.³

This manifests itself in a myriad of complaints a reader would have about the report. One, it is not clear if there is a reason for many of these variables to be in the model. Is there a good, theoretical reason why the child’s birth month should influence the relationship between school type and parent satisfaction? Does it matter if the person taking the survey is the mother or the aunt? Does it matter how many aunts or uncles are in the household? Does it matter what month the child started school? If there is no theoretical reason that these should matter, then they should not be included in the regression model. At best, extra covariates soak up additional variation in responses. At their worst, as the report does here, over 200 additional covariates undermine faith in the results of interest due to model misspecification.

Further, the report’s approach to these potential confounders, in almost every application, is to treat these confounders as fixed effects like it does with the school type variable. The school type variable makes sense as a fixed effect because it is an unordered-categorical variable (i.e., numerical values only communicate differences) with an intuitive baseline (i.e., public schools). Leaving a baseline of public schools, the fixed effects report the effect of private schools and public charters relative to a kind of “default” category. In many cases, it is not clear that this fixed effect approach generalizes across some important confounders. For example, the family dinner variable is functionally ordinal; the respondent reports that the family had anywhere from zero to seven dinners as a family. Rather than treat this variable as a single independent variable in the model, the report creates seven variables from it that all communicate the effect of, for example, zero family dinners to two family dinners, or seven family dinners to two family dinners. The baseline of two family dinners is completely arbitrary and makes little sense to the reader. The household income variable, an important confounder, is treated in a similarly confusing manner. This variable is broken into nine fixed effects with a baseline of \$0 to \$10,000 in yearly income. Rather than communicate

the effect of increasing household income, the modeling approach communicates the effect of comparing groups like the “\$150,001 or more” group or “\$30,001 to \$40,000” group to this arbitrary baseline. Even if we accept this unique way of modeling, it is an inefficient approach to handling some important variables of interest.

The oversaturation with over 230 covariates leads to some major collinearity concerns. For example, if the parent is White, the child is also likely White. The correlation between the two variables is .881 and the partial effects for both are likely to be uninformative because of multicollinearity. They are almost the same variable. The child’s birth year, grade, and school level are all effectively communicating the same thing, albeit with varying levels of precision. The availability of the internet for the respondent is going to be correlated with how many times the respondent uses the internet and, in the latter variable, it is not clear what the baseline variable is for this ordinal variable on internet consumption. It is also unclear if internet consumption should matter to the report’s central claims.

Selective Interpretation

The second shortcoming is not as obvious, but is just as important. From the first sentence, the report is making a deliberate modeling decision to isolate and emphasize the “very satisfied” responses on the parent satisfaction measures and the “strongly agree” responses for the parent’s assessment of how much the child enjoys the school. A reader versed in the statistical analysis of public opinion will find this an odd choice, especially when the item of interest is on a four-point ordinal scale from “very dissatisfied” (or “strongly disagree”) to “very satisfied” (or “strongly agree”). Almost every survey analyst will either estimate these responses on their original ordinal scale with an appropriate ordinal model or, for convenience, will collapse the items to a binary “satisfied/dissatisfied” or “agree/disagree” variable because logistic regression models are easier to communicate and estimate than ordinal models. This is standard because isolating or emphasizing a “strongly” or “very” response creates a curious baseline comparison of not, say, “satisfied/dissatisfied”, but “very satisfied/not very satisfied” because those who are “somewhat satisfied” are lumped into the group of dissatisfied respondents. This appears to be what the report does, as represented in its selective descriptive statistics in Table 1, the curious and potentially misleading model notation on p. 10, and certainly the footnotes in Tables 2 through 11. This modeling decision may potentially be masking a fuller story in the data. In essence, the early sections of the report routinely use language of “dissatisfied” to describe public school parents relative to private/charter school parents but the statistical models may not necessarily communicate that.

A Replication of the Report’s Analysis

Because the underlying data are publicly available, it is not too much effort to replicate the report’s analysis. This replication gathered the original data from the NHES and created the following variables. First, I identified the six dependent variables the report used in its analysis. I create two sets of variables that retain them on their original ordinal scale and

another set that condenses the variables into binary satisfied/dissatisfied or agree/disagree responses. I then created the school type fixed effects that the report's analysis uses.⁴

Second, I selected potential confounding effects that tried to balance model parsimony with an identification of every conceivable effect that could mitigate the relationship between school type and parent satisfaction. I include the parent's age, whether the respondent is a woman, whether the respondent has a four-year college diploma, whether the child is a woman, whether the parent is non-White, whether the child was born outside the 50 U.S. states and whether the parent was born outside the 50 U.S. states. Recall: the child's race and the parent's race are going to be highly collinear. Thus, I select the parent variable because the parent is the one taking the survey. In the case of the location of child and parent birth, those two variables are only modestly correlated ($r = .310$) and will pose no inferential problem for the statistical model. Finally, I include an interval-level measure of household income and whether the respondent received any form of public assistance.

I include two important selection variables: whether the respondent considered other schools for this child and whether they moved addresses so that the child could attend this school. Intuitively, the "consideration" variable should have a negative effect while the "moved" variable should have a positive effect because of a selection effect.

I include three family-level participation variables that work on a hypothesis that the school type variable might be sensitive to family effects and how participatory the family unit is. The first variable is the sum of whether the respondent or someone in the respondent's household attended a school event, served as a volunteer at the school, attended a school meeting, attended a parent-teacher organization meeting, attended a parent-teacher conference, participated in a school fundraiser, served on a school committee, or met with a guidance counselor. Responses range from zero (i.e., no one in the household did any of those school activities) to eight (i.e., respondent or someone in the household did all of those things since the start of the school year). I created a family activity variable that is similarly the sum of various family activity measures over the past week, including whether anyone in the family has read a story to the child, done arts and crafts with the child, played board games with the child, built something with the child, played a sport with the child, discussed time management with the child, discussed ethnic heritage with the child, visited a library with the child in the past month, visited a bookstore with the child in the past month, gone to a play with the child in the past month, visited a museum with the child in the past month, visited a zoo with the child in the past month, attended a religious event with the child in the past month, or attended a sporting event with the child in the past month. The value of this variable ranges from zero (i.e., no one in the household did any of those things with the child) to 14 (i.e., someone in the household did all of those things with the child). I also include the family dinner variable, which measures how many times the family sat down for a communal dinner involving the child in the past week. The responses intuitively range from zero to seven.

I also include two contextual factors, one for the school and another for the respondent's area. The first is a binary indicator that is a one if the total school enrollment for the child is lower than 599 students. The second is a binary indicator if the percentage of Black/Hispanic people in the respondent's ZIP code is greater than 16%.

Each statistical model has 20 covariates to explain parent satisfaction. All told, this modeling approach serves as a balance between avoiding model oversaturation while leveraging information in the survey data to identify potential sources of confounding.

First, Table 1 looks at some basic descriptive statistics for the parent satisfaction items—including the parent’s/respondent’s assessment of whether the child seems to enjoy their school experience—by school type. The values are the proportion of responses that were either “somewhat satisfied” or “very satisfied” for the first five columns or “agree” or “strongly agree” for the sixth column. The results here suggest the stark difference the report shows in the first few rows of Table 1 are a function of a decision to focus on the “very satisfied” or “strongly agree” responses.

The report’s implication from its first table is parents are more likely to be dissatisfied if their child is in public school. Instead, the stark differences look to entirely be a result of the decision to focus on just the most satisfied or most strongly in agreement. To be clear, private school and charter school parents do report higher levels of satisfaction than parents whose children are in public schools, but it would be wrong to assume that public school parents are dissatisfied with their experience. The lowest level of satisfaction for public school parents is with staff interaction. Therein, 86.8% of responses are either “somewhat satisfied” or “very satisfied.”

Table 1

School Type	Satisfied With School	Satisfied With Teachers	Satisfied With Academics	Satisfied With Order/ Discipline	Satisfied With Staff Interaction	Child Seems to Enjoy School
Charter School	0.929	0.931	0.940	0.923	0.910	0.918
Private (Catholic)	0.966	0.945	0.973	0.956	0.938	0.948
Private (Not Religious)	0.982	0.982	0.970	0.985	0.965	0.954
Private (Other Religion)	0.974	0.974	0.976	0.950	0.959	0.950
Public School	0.913	0.920	0.913	0.894	0.868	0.876

This implies two things about the main takeaways from the regression models throughout the paper. First, most of the variation between public school parents and private/charter school parents is in the relative strength or valence of their satisfaction or agreement. Second, this report seems to be a case where “statistical significance” is conflated with “substantive significance.”

Figure 1 demonstrates the first point. These are quantities of interest (i.e. simulated probabilities of a particular response) from ordinal models of the parent satisfaction items used in the report. They come from fitted ordinal logistic regressions in which every covariate is held at their typical value except for the school type variables. The “whiskers” around the point communicate 95% intervals.

Here, we see that indeed much of the variation by school type of parent satisfaction is between the “very satisfied” and “somewhat satisfied.” Differences in *dissatisfaction* by school type are negligible; indeed, more than 86% of respondents offer somewhat or very positive assessments of their experience with their child’s school. However, much of the variation in the original report might be a function of the relative valence of satisfaction or agreement. Generally: public school parents are more likely than private school or charter parents to be somewhat satisfied than very satisfied (or are more likely to agree than strongly agree) whereas private and charter school parents are more likely to be very satisfied than satisfied (or are more likely to strongly agree than just agree). This would be consistent with some of the language in the report, namely language of “report(ing) higher levels of satisfaction” for parents whose children are in private or charter school. However, this would soften some of the findings as they were originally presented.

Figure 1

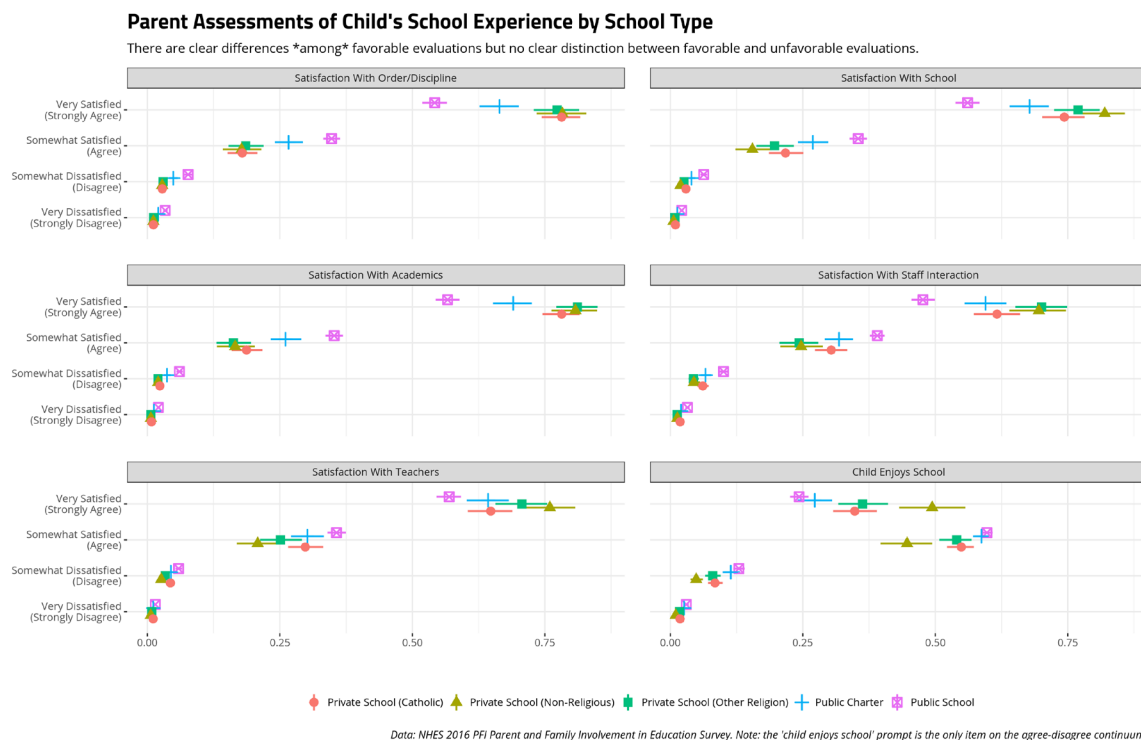
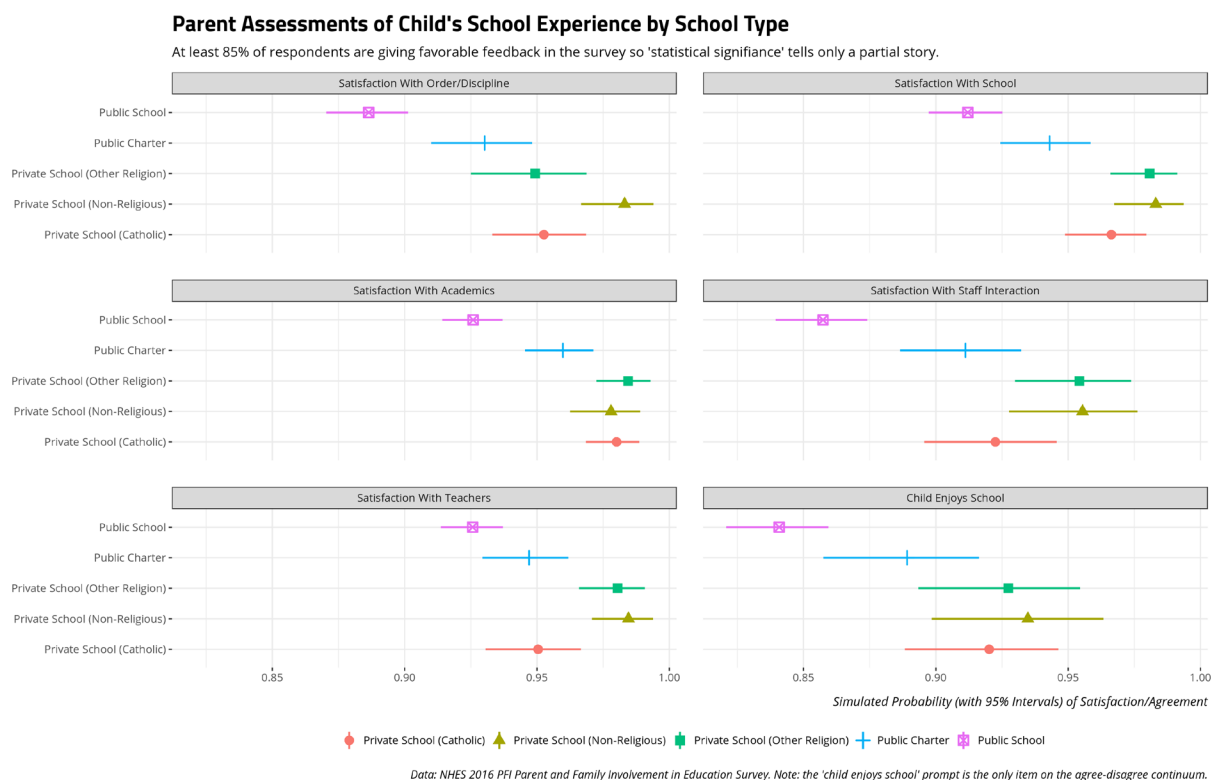


Figure 2 communicates the second point. These are quantities of interest from logistic regression models of the parent satisfaction items used in the report. These communicate the likelihood of a respondent being either satisfied or, in the case of the child enjoyment prompt, agreeing that their child seems to enjoy school. Like the results in Figure 1, they come from fitted regression models in which every covariate is held at their typical value except for the school type variables. The “whiskers” around the point communicate 95% intervals.

Figure 2



To be clear, every model summarized in Figure 1 or Figure 2 has positive and statistically significant effects for all private and charter school types relative to a baseline of a public school. This would be consistent with the report's findings. However, a reader should be careful to not conflate "statistical significance" with "substantive significance." "Statistical significance" says little of the magnitude of the effect or its importance for a reader interested in what is at stake. It only communicates whether an effect can be discerned from a counterargument of zero effect using our keenest statistical models. In this case, we see clear differences between all private school types and public schools in almost every application. There is some overlap in overall satisfaction with teachers and even differences between charter schools and public schools are discernible in just three of six models under evaluation. Notice, however, the x-axis. The underlying data suggest already high levels of satisfaction. Thus, there is a clear discernible effect of private Catholic schools relative to public schools on a parent's satisfaction with the school. However, the effect is to increase parent satisfaction from a simulated probability of satisfaction of .912 to .943. The effect amounts to an increase of three percentage points. These are significant effects, but substantively quite small.

VI. Review of the Validity of the Findings and Conclusions

The previous section suggests the following overall assessments of the validity of the findings and conclusions drawn from them. First, a reader should be skeptical of the original results provided in the report. The data source is a useful one for the hypothesis of interest, but the statistical models are misspecified and the saturation of the model with over 230 additional covariates obscures more than it illuminates. There are likewise some important misgivings about why the report chooses to emphasize the “very satisfied” and “strongly agree” responses, suggesting the results in the report are masking important heterogeneity in parent satisfaction with schools.

Second, correcting for these shortcomings in the report’s main analyses will suggest the following takeaways. First, the decision to focus on the “very satisfied” and the “strongly agree” responses mask the fact that parents are generally satisfied with their child’s school and think their child enjoys the school experience no matter the school type. That said, there is still evidence in support of the report’s original contention that charter schools and private schools increase parent satisfaction, but a decision to focus on simple regression results as proof in and of itself will miss more interesting and more important trends in the data. Namely, private schools and charter schools increase parent satisfaction with the school in as much as they increase the strength of satisfaction. Figure 1 shows private school and charter school parents are more likely than public school parents to be very satisfied than “somewhat satisfied,” while public school parents are more likely than private and charter school parents to be “somewhat satisfied” than “very satisfied.” The effect of school type does not necessarily increase satisfaction from dissatisfaction, as the report implies in the first half of the paper. This would be consistent with some insight from the cognitive dissonance literature that positive acceptance of a decision is a function of even being able to make the decision, implying here that the effect of school choice on parent satisfaction is self-fulfilling.⁵ Finally, it would be wrong to conflate “statistical significance” with “substantive significance.” The statistical significance of the school type variable only communicates whether we can discern an effect from zero, not whether the finding is substantively large or noteworthy. Indeed, Figure 2 underscores the basic descriptive statistics of Table 1 by showing that “significant” effects are so small that the reader should be wary of placing too much emphasis on the takeaways the report would like to communicate.

VII. Usefulness of the Report for Guidance of Policy and Practice

There are two superlatives of this report from a social scientific perspective. First, this is an excellent data source for explicating the hypothesis. The data set is nationally representative and sufficiently powered for the type of hypothesis testing done in the report. Second, there is reason to laud the report’s hypothesis. The normative goal of the report is clear from the first sentence, but there is at least a clear perspective and theory that informs the hypotheses under consideration.

However, the value of the report for the broader policy issue is mixed at best. The report

itself relies on model misspecification and a selective focus on just the “very satisfied” respondents to derive its fundamental takeaways. Correcting misspecification issues and generalizing the satisfaction variables will still offer some evidence consistent with the hypothesis, but it will also show additional limitations in the report’s findings. First, most of the variation among alternative school types relative to public schools is in the strength of satisfaction, rather than satisfaction levels, per se. Second, this is a clear case of confusing “statistical significance” for “substantive significance.” The effect of private and charter school types on overall satisfaction is mostly discernible, but always quite small.

Discussion about this report should focus on the following question, even if it is not clear the report offers great answers to them. First, what are the substantive effects of school types as a “treatment?” It is never enough to show regression results and treat them as sufficient. It is why social science journals are increasingly asking for “quantities of interest” like simulated probabilities and expected values to communicate some substance from the statistical models. Second, what is the underlying variation in the data? In this application, it means acknowledging that school satisfaction is already quite high. Third, and assuming there is value in looking at a “subjective” measure of school performance and that these subjective evaluations are already quite rosy for public school parents, does school choice offer a “treatment” for a problem that does not appear to exist? The report’s concluding paragraph addresses the concern of selection bias in the findings, but there is a larger point than what the report offers. School choice may not increase satisfaction as much as it increases the strength of satisfaction. School choice may offer increases in satisfaction only through the feeling of parent agency, not through anything unique or special about an alternative schooling option. Even then, it is worth reiterating that the effects are modest and concentrated mostly on those who are already at least somewhat satisfied.

This report should lead to some hesitation for those weighing some of the policy proposals from school choice advocates, especially as it pertains to curtailing the supply of public primary and secondary education through the reallocation of tax dollars. Is “school choice” a solution to a problem or a means to create the problem it purports to solve? There are important problems facing public education in the United States, but this report’s approach seems to overstate both public education’s problems and the value of the solutions it proposes.

Notes and References

- 1 For a sample of these works: Esping-Andersen, G. (1990). *The three worlds of welfare capitalism*. Princeton, NJ: Princeton University Press.

Heidenheimer, A.J. (1981). Education and social security entitlements in Europe and America. In P. Flora & A.J. Heidenheimer (Eds.), *The development of welfare states in Europe and America* (pp. 269-304). New Brunswick, NJ: Transaction Books.

Janowitz, M. (1976). *Social control of the welfare state*. New York, NY: Elsevier Scientific Publishing Co., Inc.
- 2 For a review: Bifulco R, & Ladd H.F. (2006). The impacts of charter schools on student achievement: Evidence from North Carolina. *Education Finance and Policy* 1(1), 50-90.

Booker, T.K., Gilpatric, S.M., Gronberg, T.J., Jansen, D.W. (2007). The impact of charter school student attendance on student performance. *Journal of Public Economics*, 91, 849-76.

Center for Research on Education Outcomes (CREDO). (2013). *National charter school study*. Stanford, CA: CREDO. Retrieved November 27, 2019, from https://credo.stanford.edu/sites/g/files/sbiybj6481/f/ncss_2013_final_draft.pdf

Davis, D.H., & Raymond, M.E. (2012). Choices for studying choice: Assessing charter school effectiveness using two quasi-experimental methods. *Economics of Education Review*, 31(2), 225-36.

Hanushek, E.A., Kain, J.F., Rivkin, S.G., & Branch, G.F. (2007). Charter school quality and parental decision making with school choice. *Journal of Public Economics*, 91, 823-48.
- 3 For a discussion of this from a political science perspective: Achen, C.H. (2005). Let's put garbage-can regressions and garbage-can probits where they belong. *Conflict Management and Peace Science*, 22, 327-339.

Schrodt, P.A. (2014). Seven deadly sins of contemporary quantitative political analysis. *Journal of Peace Research*, 51, 287-300.
- 4 Readers interested in these replication files should consult my Github: <https://github.com/svmiller>
- 5 For example: Festinger, L. (1957). *A theory of cognitive dissonance* Palo Alto, CA: Stanford University Press.