

Moral Convictions Often Override Concerns About Procedural Fairness: A Reply to Napier and Tyler

Linda J. Skitka · Elizabeth Mullen

Published online: 30 September 2008
© Springer Science+Business Media, LLC 2008

Abstract Napier and Tyler (this issue) question whether moral convictions about outcomes really override the influence of procedural fairness (PF) on fairness judgments and decision acceptance. The empirical answer to this question is “yes.” When people have strong moral convictions about outcomes, perceptions of outcome fairness and decision acceptance are primarily shaped by whether the morally “correct” outcomes are achieved. Pre-decision perceptions of PF have surprisingly little or no effect on these judgments. That said, pre-outcome perceptions of PF sometimes predict post-outcome perceptions of PF, even when people have morally vested outcome preferences. We provide further details supporting the validity and superiority of our data analytic approach and argue that our original conclusions were justified.

Keywords Morality · Moral mandate · Justice · Fairness · Procedural fairness · Procedural justice · Outcome fairness · Obedience to authority

Introduction

Napier and Tyler (this issue) question whether moral convictions about outcomes really override the influence of procedural fairness (PF) on fairness judgments and decision acceptance. The empirical answer to this question is “yes.” Moral convictions often have stronger effects than perceptions of PF on a variety of dependent measures. When people have strong moral convictions about outcomes, their outcome fairness judgments and decision acceptance are primarily shaped by

L. J. Skitka (✉)
Department of Psychology, University of Illinois at Chicago, m/c 285, 1007 W. Harrison St.,
Chicago, IL, USA
e-mail: lskitka@uic.edu

E. Mullen
Stanford University, Stanford, CA, USA

whether the morally “correct” outcomes are achieved. Pre-decision perceptions of PF have surprisingly little or no effect on these judgments. Do perceptions of PF nonetheless sometimes explain variance in relevant dependent measures even when people have moral mandates about outcomes? We freely acknowledge that, yes, sometimes they do. For example, Skitka (2002) found that pre-decision perceptions of PF predicted post-decision perceptions of PF (although pre-decision PF did not predict post-decision perceptions of outcome fairness or moral outrage/decision acceptance). Moreover, pre-decision perceptions of PF were stronger predictors of post-decision PF than were people’s strength of moral convictions associated with decision outcomes in that study.

Our program of research focuses on whether people’s moral convictions are psychologically important in the domain of justice judgments as well as other interpersonal contexts. This research (including the research reported in Skitka, 2002 and Skitka & Mullen, 2002) consistently supports the conclusion that when people have a moral stake in decision outcomes, their procedural and outcome fairness judgments and their willingness to comply with authorities’ decisions are influenced by the degree to which the decisions are consistent or inconsistent with their morally vested outcome preferences. By (a) suggesting that the value protection model requires that PF concerns have *no* effects; (b) selectively dismissing key dependent measures in Skitka (2002); and (c) ignoring time-order requirements of causal arguments in their re-analysis of Skitka and Mullen (2002), Napier and Tyler (this issue) make a case that there were stronger effects for PF than we claimed in these original papers. For the most part, we find these claims to be tangential to the goals of our research, that is, to test whether morally mandated outcome preferences influence perceptions of fairness. We acknowledge that pre-outcome perceptions of PF were more closely related to post-outcome perceptions of PF than people’s morally mandated outcome preferences in Skitka (2002), but stand by our conclusion that morally mandated outcome preferences had stronger effects than pre-raid perceptions of PF in Skitka and Mullen (2002). Most important to the goals of our research is not whether PF had effects, but that moral convictions did.

In summary, although morally mandated outcome preferences do not override pre-decision perceptions of PF on post-decision perceptions of PF in every instance studied, research reveals that morally mandated outcome preferences have stronger effects than pre-outcome perceptions of PF on perceptions of outcome fairness and decision acceptance in every study conducted to date. Moreover, people’s morally mandated outcome preferences often—but not always—have stronger effects than pre-outcome PF on post-outcome judgments of PF as well. Before turning to some of the more specific points that Napier and Tyler raised about our analysis and interpretation of our results, we first provide a brief review of the value protection model and its predictions, and explain how both our own and Napier and Tyler’s results are consistent with the predictions of the model.

The Value Protection Model of Justice Reasoning

The value protection model of justice reasoning proposes that moral convictions serve a number of important functions for people, including allowing them to

classify the actions of institutions, authorities, ingroup or outgroup members and even themselves into the categories of legitimate thought and deed versus a fundamental transgression. There are really two core predictions of the model—one focused on predicting people’s perceptions of decision outcomes (e.g., outcome fairness and decision acceptance) and another focused on predicting people’s perceptions of the authorities that yield these decisions (e.g., post-decision perceptions of PF, legitimacy, and related variables). To judge the degree to which the original results of Skitka (2002) and Skitka and Mullen (2002)—as well as Napier and Tyler’s re-analysis—support the value protection model, it is helpful to clearly separate these predictions, as we have in more recent versions of the model (Skitka, 2007; Skitka, Bauman, & Mullen, 2008).

The Authority Independence Hypothesis

The value protection model predicts that when people’s moral convictions are at stake, they are more likely to reason from a belief that duties and rights follow from the greater moral purpose behind rules, procedures, and authority dictates, than from the rules, procedures and authorities themselves (Skitka et al., 2008; see also Kohlberg, 1976; Rest, Narvaez, Bebeau, & Thoma, 1999). Morally based justice reasoning is not by definition anti-establishment or authority, it just is not dependent on establishment, convention, rules, or authorities. Instead, when people take a moral perspective, they focus more on their ideals, and the way they believe things “ought” or “should” be done, than on authorities. Therefore, when people have a moral stake in a decision, we predict that their fairness reasoning is less likely to be based on authorities’ and institutions’ implementation of procedures and rules, and more likely to be based on their moral intuitions about right or wrong (Haidt, 2001). Consistent with this reasoning, the *authority independence hypothesis* predicts that people’s moral beliefs are more authority independent than feelings based on preference or normative convention (see also Nucci & Turiel, 1978; Turiel, 2002). People should, therefore, feel a weaker sense of obligation to accept decisions or outcomes, or perceive decisions as fair when they believe them to be immoral, rather than simply counter to their preferences or normative expectations. In summary, the authority independence hypothesis predicts that when people have a moral stake in decision outcomes, their reasoning about outcome fairness and decision acceptance will be based more strongly on internal conceptions of personal right and wrong than on their pre-decision perceptions of authorities’ fairness or legitimacy.

In support of the authority independence hypothesis, the strength of moral conviction associated with people’s outcome preferences is an important predictor of their outcome fairness judgments (Bauman, 2007; Mullen & Nadler, 2008; Mullen & Skitka, 2006; Skitka, 2002, 2006; Skitka & Houston, 2001; Skitka & Mullen, 2002) and decision acceptance (Bauman, 2007; Mullen & Nadler, 2008; Skitka, 2006; Skitka & Mullen 2002). Moreover, consistent with the authority independence hypothesis, whether decisions are consistent or inconsistent with people’s morally convicted outcome preferences has repeatedly emerged as a stronger predictor of outcome fairness judgments and decision acceptance than pre-

decision perceptions of PF in these studies (see Skitka et al., 2008 for a more detailed review).

The Litmus Test Hypothesis

The value protection model also makes predictions about how people's personal moral beliefs will affect perceptions of authorities, or what can be usefully labeled the *litmus test hypothesis*. People often do not know the “right” answer to the kinds of decisions or conflicts authorities and institutions are asked to resolve (e.g., whether a defendant is really guilty or innocent), and therefore rely on cues that procedures are even-handed and fair as evidence that they yield good decisions. However, when people have moral certainty about what outcome authorities and institutions should deliver, they do not need to rely on procedural propriety as proxy information to judge whether the system works—in these cases, they can simply evaluate whether procedures get it “right.” “Correct” decisions indicate that leaders, authorities, and decision making procedures are appropriate and work as they should. “Wrong” answers signal that the system is somehow broken and is not working as it should. In short, *the litmus test hypothesis* predicts that people use their sense of morality as a benchmark to assess whether leaders, authorities, and procedures are valid or invalid. The litmus test hypothesis does not predict that morally mandated outcome preferences must have stronger effects on post-decision perceptions of PF than pre-decision perceptions of PF (although this point may not have been as clear as we wished in our earliest writing on this topic). Given well-known pressures toward attitudinal consistency (e.g., Festinger, 1957), pre-decision perceptions of PF are likely to have some effect on post-decision perceptions of PF. We did and do, however, clearly predict that morally mandated outcome preferences should have an effect on post-decision perceptions of PF above and beyond any effects for pre-decision PF. In support of the litmus test hypothesis, whether decisions are consistent or inconsistent with people's morally convicted outcome preferences affects people's post-decision perceptions of authorities' and institutions' PF (Bauman, 2007; Mullen & Skitka, 2006; Skitka, 2002; Skitka & Houston, 2001; Skitka & Mullen, 2002) and legitimacy (Skitka, 2006). Moreover, the effects of morally mandated outcome preferences on post-decision perceptions of PF are often as strong if not stronger than effects for pre-decision PF (Skitka & Houston, 2001; Skitka & Mullen, 2002), although we sometimes find stronger effects for pre-decision PF than for morally mandated outcome preferences on post-decision PF (e.g., Skitka, 2002).

We next discuss some of the more specific concerns raised by Napier and Tyler (this issue), and defend our analyses of our data. Specifically, we address: (a) Napier and Tyler's decision to ignore two key dependent variables examined in Skitka (2002; i.e., outcome fairness and decision acceptance), and how doing so falsely makes our conclusions sound “extreme” or misguided when they were not, (b) the approach taken in Napier and Tyler's re-analysis of Skitka and Mullen (2002), and in particular, why our tests allow stronger inferences about likely causality, (c) the problem for both the value protection model and theories of PF of failing to find interactive effects of pre-decision perceptions of PF and moral conviction on

people's post-decision perceptions of fairness, outcome acceptance, and related variables, and finally, (d) whether there is something important about issue-specific moral conviction that does not reduce to a more general individual difference variable, such as ideology or a tendency to experience all issues as morally charged.

Implications of Casting Too Narrow of a Net in the Reanalysis of Skitka (2002)

Skitka (2002) tested the value protection model by examining the effects of pre-decision perceptions of PF and morally vested outcome preferences on post-decision perceptions of outcome fairness, PF, and moral outrage (a measure that included assessments of post-decision anger as well as decision acceptance). In support of the authority independence hypothesis, our analysis (and the analysis reported in Napier and Tyler's Table 1) found that people's outcome fairness judgments and decision acceptance of hypothetical policy decisions were predicted by whether those decisions were consistent or inconsistent with perceivers' morally vested outcome preferences in each of three policy domains. Moreover, Skitka (2002) found no significant effects for pre-decision perceptions of PF on outcome fairness and moral outrage/decision acceptance measures. These results replicated across the three policy domains studied, and support the authority independence prediction that morally vested outcome preferences should emerge as stronger predictors of outcome fairness and decision acceptance than pre-decision perceptions of PF.¹ Our theory did not require that the effect for PF on these variables would be completely absent—just that the effect for PF would be weaker than the effect of morally vested outcome preferences.

Napier and Tyler (this issue), however, downplayed our results for outcome fairness and decision acceptance (i.e., tests of the authority independence hypothesis) and focused exclusively on re-analyzing our results for post-decision perceptions of PF (i.e., tests of the litmus test hypotheses). Their re-analysis of the data for post-decision perceptions of PF arrived at virtually the same conclusions as our original analyses: Pre-decision perceptions of PF were associated with post-decision perceptions of PF. More importantly to the value protection model, both analyses also found support for the litmus test hypothesis: People perceived the Supreme Court and state referenda as more procedurally fair post-decision when these procedures yielded policy decisions consistent with people's morally mandated outcome preferences, and as less procedurally fair when they did not. Support for the litmus test hypothesis also replicated across all three policy domains studied.

In short, Napier and Tyler's re-analysis did not change the core conclusions in support of the litmus test hypothesis, nor were their conclusions substantially different from our own. Our conclusions, however, of "no effects of procedures" nonetheless would seem misplaced and odd if taken out of context (i.e., without knowing the results with outcome fairness and moral outrage). That perceptions of PF would predict perceptions of PF was relatively unsurprising and was interpreted

¹ Note that Napier and Tyler (this issue) also report they found no effects of pre-threat PF on outcome fairness (see Napier and Tyler, this issue, footnote 3).

as analogous to test–retest reliability of our measure. In summary, by focusing on only one dependent variable—post-decision perceptions of PF—and ignoring post-decision perceptions of outcome fairness and moral outrage/decision acceptance—one could easily make our conclusions sound more extreme than they would in the context of all of our findings. Because pre-decision PF had no effect on perceptions of outcome fairness or moral outrage/decision acceptance, we made appropriately strong conclusions—that morally mandated ends had stronger effects than the fairness of means on people’s willingness to comply with authorities decisions and see their decisions as fair.

Re-visiting Elián: The Importance of Time–Order Relationships in Testing Causal Models

Skitka and Mullen (2002) tested the predictions of the value protection model in the context of the Elián González custody case, a long and difficult legal and public relations battle over whether a young Cuban child should be allowed to stay in the United States with extended relatives or be returned to his father in Cuba. In addition to the dramatic rescue of Elián off the Florida coast, major events in the case included a SWAT team’s forcible entry into his Miami relatives’ home to retrieve Elián, and then Elián’s eventual return to Cuba. Using a nationally representative sample and a longitudinal panel design, we assessed reactions to the case and perceptions of fairness at three junctures: several weeks before Elián was forcibly retrieved by authorities from his Miami relatives’ home; within days after his forced retrieval; and immediately after Elián was returned to Cuba (the conclusion of the case).

Testing hypotheses in the context of an engaging real world event has many advantages over testing hypotheses in the sterile confines of the lab. That said, correlational field designs make testing cause and effect relationships problematic. An important design element of both Skitka (2002) and Skitka and Mullen (2002) is that we measured people’s judgments over time, and were able to capture a clean measure of pre-decision PF—that is, one uncontaminated by knowledge about outcomes. Importantly, longitudinal panel designs provide a natural way to test hypotheses and tease apart likely cause and effect by taking into account the logical criterion that causes must occur before effects. More specifically, there are three criteria for making a causal argument: (1) changes in a proposed cause must lead to a change in the proposed effect (covariance), (2) changes in a proposed cause must occur before a proposed effect, and (3) alternative explanations must be persuasively eliminated. Survey designs that collect data at only one point in time can only establish that variables co-vary. Longitudinal panel designs can make stronger (although admittedly still imperfect) causal arguments because in addition to covariance, they allow one to establish that proposed causes occurred before proposed effects. By only testing causal models informed by time, our models were therefore more consistent with the logical requirements of testing causal hypotheses than was Napier and Tyler’s (this issue) model that used two post-decision measures to predict a third post-decision measure in their re-analyses.

Moreover, as predicted by the value protection model, people's post-decision perceptions of outcome fairness, PF, and decision acceptance converged overtime. Post-resolution judgments were all so highly correlated they may as well have been measuring a common construct (r 's between these constructs ranged from .70 to .75; we nonetheless treated them as three separate dependent measures for conceptual reasons). Any model that therefore used any one or two of the post-decision reactions to the case to predict a third (as Napier & Tyler's reanalysis of our data did) would therefore by definition fit well, but explain little because alternative models would fit just as plausibly and well.

Moreover, contrary to Napier and Tyler's (this issue) claim, we did test and compare models that included pre-raid perceptions of PF. In particular, we tested competing models that included versus excluded multiple causal pathways, according to the following conditions: (a) pre-raid measures could serve as predictors of both post-raid and post-decision measures, (b) post-raid measures could only serve as predictors of post-decision measures, and (c) post-decision measures were treated as criterion to be predicted, not predictors of other post-decision measures. In short, unlike Napier and Tyler's re-analysis, our models were explicitly guided by the time-order requirements of making a causal argument.

We also took other precautions to ensure fair tests of hypotheses. We estimated and tested all of our models using AMOS V. 4.0. As mentioned in the author notes of Skitka and Mullen (2002), a colleague independently verified all our models in LISREL, and under a wide range of assumptions. Therefore, not only did our original analyses make appropriate use of the virtues of our longitudinal panel design, but they also independently verified by another scholar using a different structural equation modeling program, to confirm the appropriateness of our approach and the soundness of our conclusions.

Because Napier and Tyler (this issue) raise questions about our model testing approach, we present more detailed results for predicting post-decision outcome fairness in Fig. 1, and Tables 1 and 2. Identical model comparisons were conducted with post-decision perceptions of decision acceptance and PF (for brevity all the additional tables are not included here, but are available from the first author). We note that these tables and figures (as well as similar ones for post-decision PF and decision acceptance) were included in the originally submitted paper for Skitka and Mullen (2002), and were therefore part of the peer-review process. However, these figures and tables were omitted from the published version of the paper due to our editor's decision that they provided too much unnecessary detail. Per the editor's request, we therefore explained that these alternative models were tested without providing extensive details about the tests, and only went into detailed descriptions of the final best-fitting model. As can be seen in Tables 1 and 2, models that included pre-decision perceptions of PF did not fit the data as well as models that omitted this latent variable when predicting post-decision perceptions of outcome fairness (note: we converged on the same best-fitting model for all three dependent measures, specifically, outcome fairness, decision acceptance, and post-decision PF). More specifically, Model A1 included all the paths depicted in Fig. 1. Model A2 dropped path d, Model A3 dropped path a, Model A4 dropped paths a and d,

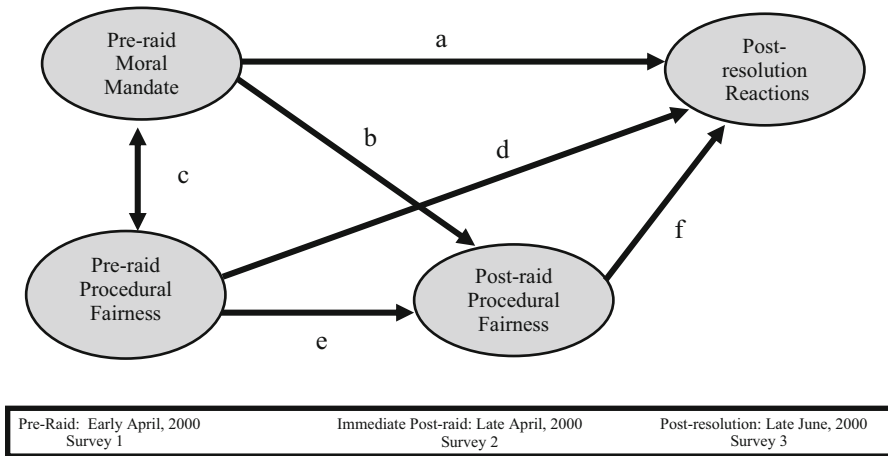


Fig. 1 The full model explaining post-resolution reactions as a function of pre-raid PF, moral mandate, and post-raid procedural fairness

Model A5 dropped paths a, b, and c, Model A6 dropped paths a, b, c, and d, Model A7 dropped paths c, d, and e, and finally Model A8 dropped paths a, c, d, and e. Each model was evaluated for improvement over the full model.

Of these, the two models that improved fit most over the full model were compared to test for “best fit”. As can be seen in Table 2, our results clearly supported the model reported in Skitka & Mullen (2002) (see also Fig. 2): Tests of competing models across all three post-resolution criteria each revealed the same conclusion: Including pre-raid judgments of PF did not yield the best fitting model; rather a model that excluded it did. More specifically, pre-raid moral convictions about how the case should be resolved directly and indirectly predicted people’s post-resolution perceptions of outcome fairness, decision acceptance, and PF. Stronger moral convictions that Elián should stay in the U.S. were associated with lower levels of decision acceptance, outcome fairness, and PF. Some of the effects of moral convictions were mediated through post-raid perceptions of PF; people’s strength of moral convictions (but not pre-raid perceptions of PF) predicted whether they felt that the raid was procedurally fair or unfair, which in turn predicted post-resolution perceptions of outcome fairness, decision acceptance, and PF.

As Napier and Tyler (this issue) note, however, we did not include direct comparisons of models that included pre-raid perceptions of PF and dropped immediate post-raid perceptions of PF in our original analyses of the Elián data (although our analyses did include tests for the direct effect of pre-raid PF on post-resolution judgments; this path was consistently nonsignificant—see models 1a, 1c, and 1e in Table 1). To test whether including post-raid PF and excluding pre-raid PF decreased model fit relative to models that included pre-raid PF and excluded post-raid PF (as Napier and Tyler argue), we conducted a number of additional model comparisons to predict post-resolution judgments of outcome fairness,

Table 1 Parameter estimates and goodness of fit indices for competing causal models for predicting perceived post-decision outcome fairness

Model parameters	Null model	Model 1a	Model 1b	Model 1c	Model 1d	Model 1e	Model 1f	Model 1g	Model 1h
<i>a</i>		.41**	.41**	–	–	–	–	.40**	–
<i>b</i>		.58**	.58**	.60**	.60**	–	–	.70**	.72**
<i>c</i>		.37**	.37**	.37**	.37**	–	–	–	–
<i>d</i>		–.04	–	–.02	–	–.02	–	–	–
<i>e</i>		.29**	.28**	.27**	.27**	.56**	.56**	–	–
<i>f</i>		.49**	.47**	.79**	.78**	.76**	.75**	.47**	.78**
χ^2	8649.83	673.17	674.52	770.23	770.39	536.74	537.02	201.78	292.06
<i>df</i>	210	159	160	161	161	96	95	64	65
NFI	.00	.92	.92	.91	.91	.92	.92	.97	.95

Note: All models that included measures of PF overtime were tested with correlated errors for each observed variable (i.e., error terms between time 1 and 2 observed indicators of PF were allowed to correlate). A chi-square of zero would indicate perfect fit; however, the chi-square statistic is known to be positively biased when sample sizes are large. Therefore, the normed fit index (Bentler & Bonett, 1980) provides a better index of model fit, where values range from 0 to 1, where 1 = a perfect fit. Models that fit above .90 indicate good fit; models that fit above .95 indicate excellent fit. Finally, all parameters except *c* are standardized regression coefficients. Parameter *c* is the covariance of moral mandate and pre-raid PF

Table 2 Model Comparisons Predicting Post-decision Perceptions of Outcome Fairness

Models	Change in χ^2	Change in <i>df</i>	Change in NFI
1a versus 1b	<i>1.35</i>	1	.00
1a versus 1c	<i>97.06*</i>	2	.01
1a versus 1d	<i>97.22*</i>	2	.01
1a versus 1e	<i>136.43*</i>	63	.00
1a versus 1f	<i>136.15*</i>	64	.00
1a versus 1 g	<i>471.39*</i>	95	.05
1a versus 1 h	<i>381.08*</i>	94	.03

Note: All changes in chi-squares, *df*, and normed fit indices represent the difference score of these values for each model as compared to the full model (1a). *Italicized* chi-squares represent a worse fit than the full model (increases in chi-square), and *nonitalicized* chi-squares represent an improved fit over the full model (decreases in chi-square). The hierarchically nested model that yields the greatest reduction in chi-square represents the best fitting model

* *p* < .01

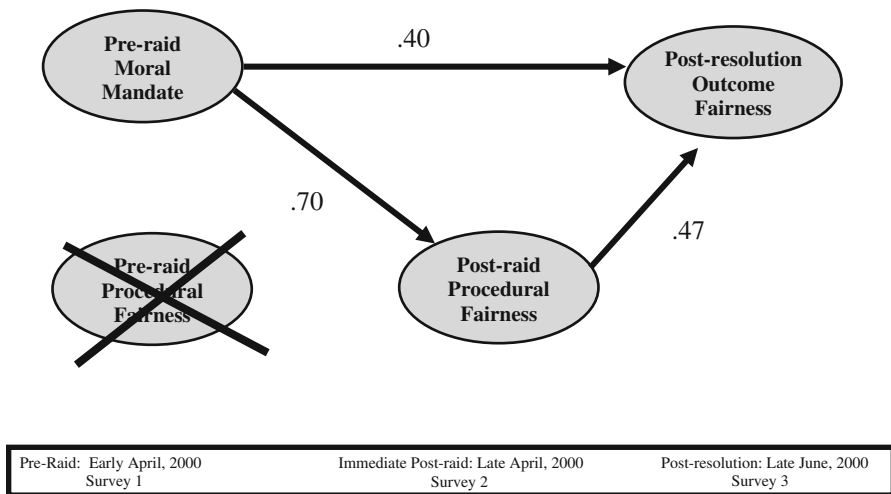


Fig. 2 The best fitting model explaining post-resolution judgments of outcome fairness

decision acceptance, and PF using AMOS 7.0 (we no longer have access to AMOS 4.0).²

Models that dropped post-raid perceptions of PF in favor of keeping pre-raid PF instead, and/or that allowed for direct effects of pre-raid PF on post-resolution judgments (with or without allowing for mediated paths through post-raid PF) fit the data significantly less well than excluding pre-raid PF from the models. The direct

² It is important to note that our models did not include outcome and procedural fairness as predictors of decision acceptance (as Napier and Tyler’s did) for the reasons noted earlier, that is, that proposed causes must occur before effects.

effect for pre-raid PF was not significant in any model that included post-raid PF, whereas post-raid PF was consistently significant regardless of whether pre-raid PF was in the model. In short, we stand by our conclusions that including pre-raid perceptions of PF not only did not improve model fit, it yielded worse fitting models than models that excluded it.

In summary, we feel our model testing approach was justified and is superior to Napier and Tyler's (this issue) for four reasons. First, unlike Napier and Tyler's model, our models only included predictors that were uncontaminated by knowledge of outcomes, and that met the important criteria for making a causal argument (i.e., proposed causes were observed and measured prior to effects). Second, because outcome fairness, decision acceptance, and post-resolution judgments of PF were all so highly correlated post-decision, it was inappropriate to claim that one post-resolution variable was causal of the other. Thus, contrary to Napier and Tyler's model that included paths from post-resolution outcome fairness and PF to predict decision acceptance, our models did not include paths from post-resolution variables to predict other post-resolution variables. Third, our original analyses as well as new analysis in light of Napier and Tyler's comments indicated that models that excluded pre-raid perceptions of PF fit the data better than models that excluded post-raid perceptions of PF instead. Finally, one of the criteria for a good theory is parsimony. Our best fitting model was more parsimonious than the model tested and advocated by Napier and Tyler. Taken together, we feel our model testing approach and conclusions were appropriate and valid, and represent more appropriate and valid tests of hypotheses than the approach used by Napier and Tyler.

Finally, our final models make sense in light of the dramatic events associated with Elián's case. It seems highly improbable that people's reactions to the resolution of Elián's case would be virtually unaffected by media images of armed SWAT teams retrieving a 5-year-old boy by force from his relatives' home. Our open-ended reactions to the case (also reported in Skitka & Mullen, 2002) found considerable evidence of strong reactions to the raid, ranging from "about time!" to strong condemnation of governmental officials and their agents. More open to question was not whether post-raid reactions would influence post-resolution judgments, but whether post-raid reactions would be predicted by people's moral convictions about how the case should be resolved, their pre-raid perceptions that authorities were procedurally fair, or both. Our data supported the first of these possibilities better than the other two.

Implications of an Elusive Moral Mandate by PF Interaction

Napier and Tyler (this issue) argue that the absence of a PF by MM interaction on perceptions of fairness is problematic for the value protection model. In particular, Napier and Tyler argue that the value protection model predicts that when people have a moral stake in the decision outcome that there should be no influence of procedures on people's perceptions of outcome fairness and decision acceptance. We reject this interpretation of the value protection model and instead argue that

when people have a moral stake in the outcome their perceptions of outcome fairness and decision acceptance will be better predicted by whether outcomes are consistent with their moral mandate than by PF. In short, the appropriate comparison is strength of effects, not presence or absence of effects when people have moral mandates. We further argue that when people do not have a morally mandated outcome preference, the value protection model predicts that positive aspects of procedures should enhance perceptions of outcome fairness and decision acceptance, as found in research documenting fair process effects (e.g., Folger, 1977; Van den Bos, Wilke, & Lind, 1998). These predictions imply that one would expect a PF by moral mandate interaction, that is, that one should observe stronger effects for PF when people do not have moral mandates about outcomes than when they do.

Napier and Tyler (this issue) claim that because there was no evidence of a PF by moral mandate interaction for perceptions of outcome fairness and moral outrage in Skitka (2002) that the pattern of the data necessarily looks like the right panel of their Fig. 1. However, our results as well as Napier and Tyler's reanalysis, instead indicated that the slopes for PF on perceptions of outcome fairness and moral outrage are all flat (irrespective of whether participants had a moral mandate, or the type of mandate) because PF did not significantly influence people's perceptions of outcome fairness or moral outrage/decision acceptance. In short, the absence of a moral mandate by PF interaction was not due to a consistent positive effect of PF on outcome fairness irrespective of moral mandate (as depicted in the right panel of their Fig. 1); instead, there was *no* significant effect for PF on perceptions of outcome fairness and moral outrage/decision acceptance, regardless of level of moral mandate. Clearly, the absence of interactive effects for PF and MM is surprising: both the value protection model and theories of PF would predict at a minimum that one should see stronger effects for PF when people do not have strong moral convictions about outcomes. Correlational field designs, however, cannot provide definitive tests of whether these variables interact because of well-known problems with having sufficient observations at the joint extremes of any two given variables (McClelland & Judd, 1993). Adding a third variable—whether perceivers supported or opposed a decision (as Napier and Tyler did)—places an even greater burden on statistical sensitivity. For example, one would need large numbers of people who thought the procedures used to decide that the Elián case were maximally procedurally fair or unfair, and *within* those extremes, one would also need people who were maximally and minimally morally convicted about the outcome of the case. On top of these already stringent requirements for statistical power, one would also need approximately equal number of these people to be for or against Elián's return to Cuba within each of the "cells" of the PF (maximally high, low) by moral conviction (maximally high, low) joint distributions. Given that responses to these variables were normally distributed, with most of the observations on any given variable in the middle rather than the extremes of their distributions, the power to detect interactions is *vanishingly* small. The "four-corners" problem is not unique to our studies or methods, but is common to virtually all field designs that do not explicitly sample cases from the joint extremes of variables of interest (McClelland & Judd, 1993).

Consistent with the notion that the nonsignificant interactive effects observed in Skitka (2002), Skitka and Mullen (2002) and Napier and Tyler (this issue) are a consequence of the low statistical power to detect these effects in correlational field designs is the fact that we more often observe the predicted effects of PF for nonmandated participants when we study moral conviction using experimental methods, or specifically select research participants as a function of whether they do or do not have moral mandates about outcomes (see Bauman & Skitka, *in press*, Study 2; Skitka & Houston, 2001). That said, a number of these studies have still found no evidence for a moral mandate by PF interaction on perceptions of outcome fairness or decision acceptance even when manipulations of PF had strong effects on PF manipulation checks (e.g., Bauman & Skitka, *in press*; Mullen & Nadler, 2008). Taken together, the failure to find interactive effects would seem to be more problematic for theories of PF than the value protection model. We find the predicted effects for moral convictions; the predicted effects for PF are much more elusive.

That said, more research using sufficiently powerful designs to allow for the detection of interaction effects is needed. Given that a great deal of our research—even using designs better suited for detecting interactive effects—nonetheless finds weak or no effects for PF on perceptions of outcome fairness and decision acceptance, a number of other explanations for the lack of PF effects in moralized contexts should be explored in future research as well. For example, perhaps there is something special about decisions in moral domains that focuses people’s attention more on outcomes than procedures, regardless of whether they have a personal moral investment in these decisions. According to theories of moral politics (e.g., Mooney, 2001), all it takes to moralize an issue is for one side to frame their position in moral terms. It is one thing to attempt to resolve conflict over competing preferences; it is quite another to resolve conflict when one (or more) factions frame the issue in terms of first order truth (e.g., Skitka, Bauman, & Sargis, 2005, Study 4). To support alternatives or possible compromises to what one side sees as “right,” “moral,” and “good” is to be absolutely “wrong,” or “immoral,” if not “evil” (e.g., Black, 1994; Bowers, 1984; Meier, 1994; Skitka et al., 2005; Tetlock, Kristel, Elson, Green, & Lerner, 2000), even if the compromiser has no personal moral investment in the decision at hand. Thus, even people without a moral stake in a given decision may well recognize that others see the issue in moral terms, which may in turn change the relative importance they would otherwise place on PF. Theories of PF would seem to predict that these would be the very circumstances that should make people especially likely to turn to procedures and the imprimatur of the state to endorse one or another set of values. The longstanding debate about the morality of abortion despite the 1974 Supreme Court decision in *Roe v. Wade*, however, reveals that the imprimatur of the state is not always successful at resolving conflicts about moral politics (Hunt, 1999).

In summary, we do not see the lack of observed interaction effects in Skitka (2002) and Skitka and Mullen (2002) as a fatal blow to either the value protection model or theories of PF, however, it may represent a more serious problem for theories of PF than the value protection model. The low statistical power common to correlational field designs is a very plausible explanation for failing to detect the

predicted interaction effect (albeit perhaps the least interesting explanation). Worth exploring in future research are alternative explanations to low statistical power, such as whether people become more focused on outcomes than procedures when they recognize that others see issues in terms of first order moral principles, even when they themselves do not.

Is There Something Special about Issue-Specific Moral Conviction?

Napier and Tyler (this issue) questioned whether the moral mandate construct captures a more generalized moral worldview or ideology, rather than issue-specific moral conviction attached to a single issue domain. The lack of perfect coherence and constraint in people's political belief systems is a well-known limitation of using ideology to understand who will feel strongly about certain issues and why (e.g., Converse, 1964). Similarly problematic are the generally weak ties between people's commitments to abstract ideological beliefs and values, and their attitudes about specific issues of the day. For example, some people claim that they deeply value the sanctity of life and see their commitment to this value as central to their ideological belief system. If a commitment to this value serves as the central organizing theme of their ideological beliefs, studies should find strong correlations between a variety of life-relevant policy positions. However, studies consistently find weak correlations between opposition to abortion and a more generalized pro-life stance on issues such as capital punishment or voluntary euthanasia (e.g., Darwin, 1982; Lester, 2000).

One reason why there are often gaps between commitments to abstract universal values and individual expressions of moral conviction may be because there are so many universal abstract moral values (e.g., justice, loyalty, and liberty) that they conflict and cannot be simultaneously maximized. The emotions people experience in connection with moral convictions may help signal which values are most important to maximize in specific contexts (cf. Haidt, 2001). Taken together, although one might see some ideological coherence and consistency in the degree to which people express moral convictions across issues, we posit that there will be important increases in explained variance in people's reactions to decisions if one takes into account the degree to which people have a strong moral stake in the particular decision at hand.

Consistent with the position that there is something special about issue-specific moral conviction, some of our other research has demonstrated that strength of moral conviction does not reduce to variables such as political orientation and strength of political orientation, party identification and strength of party identification (e.g., Skitka & Bauman, 2008) or a wide variety of nonmoral characteristics of attitude strength (e.g., extremity, importance, certainty, centrality; Skitka et al., 2005). Moreover, issue-specific moral conviction explains unique variance in dependent measures after controlling for people's strength of moral conviction associated with several unrelated variables (Skitka et al., 2005, Study 2). Although moral convictions on unrelated issues often explain unique variance in reactions to specific issues, issue-specific moral conviction (as predicted) nonetheless continues to explain unique variance as well. Therefore, although ideology or a

generalized tendency to see issues in a moral light does explain variance in the settings we have studied to date, knowing whether a person's specific position on a given issue is held with strong moral conviction explains significant additional variance even when controlling for these possible alternative explanations for moral mandate effects.

Napier and Tyler's (this issue) analysis presented in Table 4 is consistent with our conclusion that issue-specific moral conviction explains significant unique variance in relevant dependent measures, in this case, moral outrage. Even when one statistically controls for "generalized moral view" by including moral convictions about unrelated issues into the prediction equations, issue-specific moral conviction associated with immigration, abortion, and civil rights consistently explained more unique variance in moral outrage about those respective issues than did moral convictions about nonrelevant issues. We conducted additional analysis of the Skitka (2002) data to explore how much unique variance issue-specific moral conviction explained in outcome and post-decision PF as well. Controlling for pre-decision PF and moral conviction about unrelated issues, issue-specific moral convictions explained 15%, 45%, and 39% unique variance (R^2 change) in perceptions of outcome fairness across the respective issues of civil rights for homosexuals, abortion, and immigration (all $p < .01$). Issue-specific moral conviction explained 3%, 11%, and 6% variance in post-decision perceptions of PF across the same issues and using the same controls (all $p < .01$). Therefore, even very conservative tests of the relative power of issue-specific moral convictions support the value protection model and our broader integrated theory of moral conviction that there is something privileged about issue-specific moral conviction that does not reduce to a generalized moral orientation or worldview (see also Skitka & Bauman, 2008; Skitka et al., 2005).

Conclusions

Taken together, research using a variety of different measures has converged on support for the basic conclusion that there are important differences as a function of whether feelings about outcomes or attitude objects are experienced as preferences or moral convictions. That said, there is no end (at least in principle) to the potential for hair-splitting over the degree to which one has created a level methodological or statistical playing field for testing hypotheses. To address this concern, the moral mandate program of research has tested hypotheses using a wide range of operationalization of variables, contexts, and methods (longitudinal field designs, experiments, archival analysis, etc.) and consistently converged on the same conclusions: When people have strong moral convictions about outcomes they often become more focused on whether those outcomes are achieved than how they are achieved. For example, one experiment found that vigilantism is perceived to be as procedurally fair as due process when it achieves morally mandated ends (Skitka & Houston, 2001). Do these conclusions mean that PF never matters? Certainly not. Although the effects for pre-decision PF do not replicate as well across studies as the effects we have observed for morally vested outcome preferences, we

nonetheless do find in some circumstances that pre-decision PF has stronger effects on post-decision PF than do people's morally convicted outcome preferences (e.g., Skitka, 2002).

Taken together, our approach assumes that not everyone looks at a given situation with the same motivational priorities. Consider the example of a university considering whether to use student fees to fund abortions at a student health clinic (e.g., Bauman, 2007). One perceiver might be concerned primarily about whether her student fees are likely to increase and how this will affect her budget. Because this student is taking an economic perspective on this issue, the criteria she uses to decide whether using student fees to fund abortions are fair may be quite different from a second student, who sees the issue more in terms of his social identity as a member of the university. This second student may be more concerned about whether he has voice in the decision making process about whether to increase student fees, and whether authorities are respecting students' opinions. A third student may perceive the issue in terms of the fundamental morality of blocking or facilitating students' access to abortion services. All three students are viewing exactly the same reality, but from very different points of view. What they see as at stake in the decision—their economic self-interest, their status and standing, or basic questions of right or wrong—has considerable potential to influence whether the decision-making procedures and the decision outcome are perceived as fair.

A theory of justice that takes multiple motivational perspectives into account and develops contingent predictions about how people think about fairness within each of these frames of reference has considerable promise to integrate when people will weight different concerns more heavily into their justice reasoning, and therefore the relative weight they place on outcomes versus procedures on their perceptions of fairness (see Heuer, Penrod, Hafer, & Cohn, 2002; Schroeder, Steel, Woodrell, & Bembenek, 2003; Skitka, 2003; Skitka, Aramovich, Lytle, & Sargis, *in press*, for recent efforts in this direction). Although much of our research to date has focused on testing a theory that predicts that morally mandated outcome preferences have an effect on justice judgments and related variables, we argue that it is important to move toward research that tests predictions from integrative theories that make contingent predictions about how different motivational priorities affect how people perceive the fairness of the same situation. A motivationally contingent theory of justice can account for the mundane reality that people often disagree about whether a given situation was handled fairly or unfairly. Moreover, the notion that people are likely to approach the same situation from different perspectives, that in turn shapes the fairness norms or considerations they apply to it, suggests that future research should extend beyond the study of how individuals in isolation make fairness judgments. Future research should begin to explore how people socially negotiate and arrive at consensus about fairness.

Acknowledgments The research originally described in Skitka (2002) and Skitka and Mullen (2002) and the focus of this reply were supported by grants from the National Science Foundation (#9617861, #0084781). We thank Christopher Bauman, Brad Lytle, Scott Morgan, Sylvia Perry, Benjamin Saunders, and Daniel Wisneski for sharing their reactions to Napier and Tyler's (this issue) critique, and especially Christopher Bauman and Brad Lytle for their comments on an earlier draft of this article.

References

- Bauman, C. (2007). Procedural and moral influences on fairness judgments, decision acceptance, and organizational behavior (Doctoral dissertation, University of Illinois at Chicago, 2006). Dissertation Abstracts International, 67, 5463.
- Bauman, C. W., & Skitka, L. J. (in press). Moral conflict and procedural justice: Moral mandates as constraints to voice effects. *Australian Journal of Psychology*.
- Bentler, P. M., & Bonett, D. G. (1980). Significance tests and goodness of fit in the analysis of covariance structures. *Psychological Bulletin*, 88, 588–606.
- Black, G. D. (1994). *Hollywood censored: Morality codes, Catholics, and the movies*. Cambridge, U.K.: Cambridge University Press.
- Bowers, W. J. (1984). *Legal homicide: Death as punishment in America, 1864–1982*. Boston: Northeastern University Press.
- Converse, P. E. (1964). The nature of belief systems in mass publics. In D. Apter (Ed.), *Ideology and discontent* (pp. 201–261). New York: Free Press.
- Darwin, D. O. (1982). Public attitudes toward life and death. *Public Opinion Quarterly*, 46, 521–533.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Folger, R. (1977). Distributive and procedural justice: Combined impact of “voice” and improvement of experienced inequity. *Journal of Personality and Social Psychology*, 35, 108–119.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814–834.
- Heuer, L., Penrod, S., Hafer, C. L., & Cohn, I. (2002). The role of resource and relational concerns for procedural justice. *Personality and Social Psychology Bulletin*, 28, 1468–1482.
- Hunt, A. (1999). *Governing morals: A social history of moral regulation*. Cambridge, UK: Cambridge University Press.
- Kohlberg, L. (1976). Moral stages and moralization: The cognitive developmental approach. In T. Lickona (Ed.), *Moral development and behavior: Theory, research and social issues*. New York: Holt, Rinehart, and Winston.
- Lester, D. (2000). Is there a consistent pro-life attitude? *Advances in Psychology*, 1, 1–9.
- McClelland, G. H., & Judd, C. M. (1993). Statistical difficulties of detecting interactions and moderator effects. *Psychological Bulletin*, 114, 376–390.
- Meier, K. J. (1994). *The politics of sin*. Armonk, NY: M. E. Sharpe.
- Mooney, C. Z. (2001). The public clash of private values: The politics of morality policy. In C. Z. Mooney (Ed.), *The public clash of private values: The politics of morality policy* (pp. 3–20). New York, NY: Seven Bridges Press.
- Mullen, E., & Nadler, J. (2008). Moral spillovers: The effects of moral violations on deviant behavior. *Journal of Experimental Social Psychology*, doi: [10.1016/j.jesp.2008.04.001](https://doi.org/10.1016/j.jesp.2008.04.001).
- Mullen, E., & Skitka, L. J. (2006). Exploring the psychological underpinnings of the moral mandate effect: Motivated reasoning, identification, or affect? *Journal of Personality and Social Psychology*, 90, 629–643.
- Napier, J. L. & Tyler, T. R. (2008). Does moral conviction really override concerns about procedural justice? A reexamination of the value protection model. *Social Justice Research*, doi: [10.1007/s11211-008-0083-y](https://doi.org/10.1007/s11211-008-0083-y)
- Nucci, L. P., & Turiel, E. (1978). Social interactions and the development of social concepts in pre-school children. *Child Development*, 49, 400–407.
- Rest, J. R., Narvaez, D., Bebeau, M. J., & Thoma, S. J. (1999). *Postconventional moral thinking. A neo-Kohlbergian approach*. Mahwah, NJ: Lawrence Erlbaum.
- Schroeder, D. A., Steel, J. E., Woodrell, A. J., & Bembek, A. F. (2003). Justice within social dilemmas. *Personality and Social Psychology Review*, 7, 394–387.
- Skitka, L. J. (2002). Understanding judgments of fairness in a real-world political context: A test of the value protection model of justice reasoning. *Personality and Social Psychology Bulletin*, 28, 1419–1429.
- Skitka, L. J. (2003). Of different minds: An accessible identity approach to why and how people think about justice. *Personality and Social Psychology Review*, 7, 286–297.
- Skitka, L. J. (2006). *Legislating morality: How deep is the U.S. Supreme Court's reservoir of good will?* Paper presented at the meeting of the International Society for Justice Research, Berlin, Germany.

- Skitka, L. J. (2007). *Morality as a constraint on deference to authority*. Paper presented at the 11th Ontario Symposium: Justice and Legitimacy, Waterloo.
- Skitka, L. J., Aramovich, N. P., Lytle, B. L., & Sargis, E. G. (in press). Knitting together an elephant: An integrative approach to understanding the psychology of justice reasoning. In R. Bobocel, A. Kay, J. Olson, & M. Zanna (Eds.), *The psychology of justice and legitimacy, 11th Ontario Symposium*.
- Skitka, L. J., & Bauman, C. W. (2008). Moral conviction and political engagement. *Political Psychology, 29*, 29–54.
- Skitka, L. J., Bauman, C. W., & Mullen, E. (2008). Morality and justice: An expanded theoretical perspective and review. In K. A. Hedgvedt & J. Clay-Warner (Eds.), *Advances in group processes*. Elsevier: Amsterdam.
- Skitka, L. J., Bauman, C. W., & Sargis, E. G. (2005). Moral conviction: Another contributor to attitude strength or something more? *Journal of Personality and Social Psychology, 88*, 895–917.
- Skitka, L. J., & Houston, D. (2001). When due process is of no consequence: Moral mandates and presumed defendant guilt or innocence. *Social Justice Research, 14*, 305–326.
- Skitka, L. J., & Mullen, E. (2002). Understanding judgments of fairness in a real-world political context: A test of the value protection model of justice reasoning. *Personality and Social Psychology Bulletin, 28*, 1419–1429.
- Tetlock, P. E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology, 78*, 853–870.
- Turiel, E. (2002). *The culture of morality: Social development, context, and conflict*. Cambridge, UK: Cambridge University Press.
- Van den Bos, K., Wilke, H. A. M., & Lind, E. A. (1998). Evaluating outcomes by the fair process effect: Evidence for different processes in fairness and satisfaction judgments. *Journal of Personality and Social Psychology, 74*, 1493–1503.