



Munich Personal RePEc Archive

The Use of Identity Primes to Explain Behavioral Differences Between Groups: A Methodological Note

van Hoorn, Andre

May 2017

Online at <https://mpra.ub.uni-muenchen.de/80011/>
MPRA Paper No. 80011, posted 06 Jul 2017 05:09 UTC

The Use of Identity Primes to Explain Behavioral Differences

Between Groups: A Methodological Note

André van Hoorn

Radboud University, Institute for Management Research, PO Box 9108, 6500 HK, Nijmegen, the Netherlands. Tel: +31 243 611 663, Fax: +31 243 612 379, E-mail: A.van.Hoorn@fm.ru.nl.

The Use of Identity Primes to Explain Behavioral Differences Between Groups: A Methodological Note

Abstract

Economists are increasingly using primes that make group identity salient to overcome the inferential limitations of behavioral quasi-experiments involving pre-existing groups (e.g., males vs. females). However, while priming group identity provides powerful means for identifying a causal effect of group membership on individuals' preferences, so far, there has been little methodological reflection on the use of identity primes to identify the causes of group differences in preferences. This note's main contribution is to offer a framework for thinking systematically about the treatment effects of priming individuals' group identity and the identification of specific group traits explaining differences in preferences or behavior between pre-existing groups. The framework sets a high bar for studying the causes of group differences in preferences using identity primes but we clarify its usefulness using concrete examples.

Keywords: Experimentation; random assignment; salience; quasi-experiment; group membership; culture

JEL codes: C90; Z10; C36

1. Pre-Existing Groups and Experiments as Measurement Tools

This methodological note provides a systematic discussion of the conditions under which researchers are able to offer clean identification of the causes of differences in behavior or preferences between pre-existing groups, for instance, males vs. females (Croson and Buchan 1999; Eckel and Grossman 2008; Gneezy et al. 2009; Gong et al. 2015; Visser and Roelofs 2011) or Western vs. non-Western cultures (Cason et al. 2002; Gächter et al. 2010; Henrich et al. 2001; Oosterbeek et al. 2004; Roth et al. 1991). A large literature exists that insightfully applies experimental games and other choice situations to study a range of economic issues. Croson and Gächter (2010) list various ways in which experiments contribute to economic research by testing theoretical predictions, by providing abstract behavioral models and by helping the refinement of existing theories and the construction of new theories. In addition, experiments can act as a measurement tool, meaning that experiments provide a rigorous means for measuring individuals' preferences and, particularly, group differences therein (see, also, Guala 2008 and Plott and Smith 2008, p. xi). Camerer and Fehr (2004) provide detailed guidelines for the use of experimental economic games to measure and compare preferences across different human groups identified beforehand.

As illustrated by a substantial body of research, different games can be used—for example dictator games, trust games or coin tossing games—and the main challenge for researchers is to create matched samples for each distinct group considered. If selected samples do not diverge on the most important (economic) parameters, any difference in choice behavior observed for the groups demonstrates group differences in preferences (of course, we equate preferences with behavior and vice versa). In practice, researchers most often aggregate experimental behavior and compare average scores across groups to examine the extent to which preferences differ between these groups. Alternatively, data on experimentally elicited behaviors can be used to estimate a generic utility function but with specific parameters that are allowed to vary across groups (see, for example, Andersen et al. 2008; Bellemare et al. 2008; Von Gaudecker et al. 2011).

An important distinction between the use of experiments as a measurement tool or as a means for

testing, refining and constructing economic models and theories is that the nature of the scientific inference substantiated by the results of the experiments is different. In the latter case, there would typically be random assignment and therefore clean identification of causal effects. In the former case, however, the experimental results come in the form of measured differences between pre-existing groups, say American students and Chinese students, but without identification of the causes underlying any observed differences between the groups. The classic difference is between true experiments and quasi-experiments, where the former requires random assignment to treatment, which is not achieved when groups are pre-existing and one simply compares average (game) behavior in one group with average (game) behavior in another group (Shadish et al. 2002). As an example, consider an experiment on Sino-American differences in preferences. In this case, identifying the causal effect of being American or Chinese on individuals' preferences would require the researcher to assign subjects randomly to be either American or Chinese, which is not possible (Berry 2002; Burnham and Kurzban 2005; Van de Vijver and Leung 1997; Van Hoorn 2012). However, without such random assignment the evidentiary basis for claiming that any observed differences are causally due to an individual being American or Chinese remains as strong (or weak) as it would be in any quasi-experimental research design (Burnham and Kurzban 2005; Campbell and Cook 1979).

2. From Measuring Group Differences to Explaining Group Differences

Studies that use behavioral experiments as a measurement tool typically provide some sort of quantification of observed behavioral differences between pre-existing groups, which allows them to address the question whether and to what degree the groups considered exhibit different preferences. The majority of the literature on gender or cross-cultural differences in preferences follow this format (Cason et al. 2002; Croson and Buchan 1999; Eckel and Grossman 2008; Gächter et al. 2010; Gneezy et al. 2009; Gong et al. 2015; Henrich et al. 2001; Oosterbeek et al. 2004; Roth et al. 1991; Visser and Roelofs 2011). The literature is not limited to considering these specific groups, however, and several other interesting social groups have been considered as well, including children (Harbaugh et al. 2002), elderly (Charness

and Villeval 2009), economics students (Marwell and Ames 1981), finance professionals (Haigh and List 2005), car mechanics (Beck et al. 2014) and CEOs (Fehr and List 2004). More importantly, the inferential limitations of such quasi-experimental evidence on group differences that we discussed above are increasingly well recognized. Several studies have thus worked on moving beyond measurement to provide a causal account of experimentally observed behavioral differences or differences in preferences between groups.

Following a large literature in cross-cultural psychology (Bond and Van de Vijver 2011; Leung and Van de Vijver 2008; Matsumoto and Yoo 2006; Poortinga 2016), two main approaches for strengthening the inferential power of experiments involving pre-existing groups and identifying the causes of group differences in preferences can be discerned.¹ The first is the use of primes to make group identity salient among the members of the specific groups considered, what can be called the identity-priming or “systematic contrast” approach (Leung and Van de Vijver 2008, p. 154). Priming group identity allows researchers some degree of control over group membership, providing a means to create randomization and realize a research design that is closer to truly experimental (Shadish et al. 2002) even in situations where group membership itself is not randomized. When using identity primes, behavioral differences between group members that have received the identity prime and group members that have not received the identity prime can be validly interpreted as reflecting the causal effect of belonging to this specific group on individuals’ preferences. Benjamin et al. (2010) is an example of a study taking the identity-priming approach, finding that priming Asian identity of Asian-Americans had a positive effect on subjects’ time preferences. Similarly, Benjamin et al. (2016) primed individuals belonging to different religious denominations with their religious identity, finding that the prime affected preferences in some domains, for instance, risk aversion but not others. Two other examples are the studies by Cohn et al.

¹ Note that cross-cultural psychology has a long-standing interest in quantifying and explaining group differences in such constructs as attitudes or value orientations, which resonates with economists’ interest in analyzing group differences in preferences.

(2014) and Cohn et al. (2015). The former study primed bankers with their professional identity, finding that this prime increased cheating in a coin tossing game. The latter study primed the criminal identity of inmates, also finding a positive effect on cheating in a coin tossing game.

The second approach for strengthening the inferential power of experiments involving pre-existing groups involves the use of additional data or evidence to provide further analysis of the experimentally observed group differences and argue a specific (causal) interpretation of these differences. This approach can be called the “co-variate” approach (Leung and Van de Vijver 2008, p. 154). Assuming one’s sample comprises a sufficiently large number of groups, one can use regression analysis to provide some evidence on the reason why different groups, on average, behave differently. Herrmann et al. (2008) is an example of a study taking the co-variate approach, presenting correlations between cross-country differences in anti-social punishment and various indicators of national culture. In similar fashion, Gächter and Schulz (2016) find that cross-country differences in dishonest game behavior correlate with the historical prevalence of political fraud, tax evasion and corruption in these countries. Another example is the study by Voors et al. (2012), which finds that differences in experimentally elicited risk, time and altruistic preferences among regions in Burundi correlate with differences in the history of violent conflict in these regions. Finally, Apicella et al. (2014) find that living in an isolated region or not predicts differences in the strength of the experimentally observed endowment effect among different groups of Hadza Bushmen in Northern Tanzania.

3. Omitted Variables and Identification of the Causes of Group Differences

Both the identity-priming approach and the co-variate approach provide researchers with interesting tools for pushing the causal inferences substantiated by experiments involving pre-existing groups. At the same time, however, these two approaches are not a panacea, as their ability to offer causal understanding of the factors driving behavioral differences between groups can still be limited (Leung and Van de Vijver 2008). In the first part of this section we elaborate the major threat to clean identification of the causes of differences in preferences in both co-variate studies and identity-priming studies, which boils down to a

potential omitted variable bias. As the potential for omitted variable bias is most obvious in co-variate studies, I consider the problem in the context of this type of studies first, before elaborating the potential relevance of omitted variables in identity-priming studies. The second part of this section draws on ideas behind instrumental variable analysis to offer a framework for thinking systematically about omitted variables bias in identity-priming studies and the use of identity primes to identify the causes of differences in preferences between pre-existing groups.

3.1. Omitted variables and the sources of behavioral differences between groups

As indicated, the co-variate approach involves taking data on preferences rigorously measured using experiments and relating these data to other data (Leung and Van de Vijver 2008). In practice, the co-variate approach tends to involve estimating a regression model with experimentally observed behavior as the dependent variable and other measures, for example national culture indicators (e.g., Herrmann et al. 2008), as independent variables. The problem of omitted variable bias subsequently occurs naturally in this situation in the same way that it does in any observational study. There are many possible factors that, in principle at least, can account for systematic differences in experimentally observed behavior and it is not feasible to specify a regression model that does not allow for alternative explanations of the observed variance. Thus, for instance, national culture may be found to correlate with country differences in preferences but based on the results of the regression analysis alone we are not able to rule out that a factor other than culture is driving this relationship. In many cases, the ability of a co-variate study of group differences in preferences to reduce the problem of omitted variables is also limited because the number of groups considered is relatively small (Matsumoto and Yoo 2006). The study by Herrmann et al. (2008), for instance, considers individuals from 16 countries while the study by Apicella et al. (2014) compares eight different groups of Hadza Bushmen.

Whereas the concern for an omitted variable bias is relatively obvious when using the co-variate approach to studying group differences in preferences, this concern is more subtle when using the identity-priming approach. The rationale for applying an identity prime is to create some randomness

concerning individuals' membership to specific human groups, even among individuals that belong to the same pre-existing group. The identity prime is thereby intended to make group membership more salient compared to a control group that has not received the identity prime. At the same time, the prime itself is randomly assigned, meaning that also the salience of group membership is random across individuals (even though, as indicated, strictly speaking, the individuals still belong to the same human group). Used in this way, the identity-priming approach has no difficulty offering clean causal identification of the effect that group membership can have on individuals' preferences. The problem of omitted variable bias becomes relevant, however, when taking the next step, which is to identify the causes of group differences in preferences, i.e., to identify the specific group-level traits that explain why certain groups are, say, more risk averse than other groups are. Identity primes work to make group membership salient in a way that can affect behavior. It is not clear, however, exactly which feature of individuals' groups is made salient by this prime (Leung and Van de Vijver 2008). An omitted variable bias can thus occur because priming individuals' group identity can make multiple distinct group-level traits salient at the same time, all of which could account for a found treatment effect (Schwarz and Strack 1981). Hence, without additional information it is not possible to identify exactly what it is about groups or group membership that makes that individuals belonging to some human groups behave so differently from individuals belonging to other human groups, e.g., men vs. women or Americans vs. Chinese.

3.2. Identification of the causes of group differences in preferences

The above discussion introduced the problem of omitted variable bias as a threat to the ability of co-variate and identity-priming studies to identify the causes of group differences in preferences. A logical next step is to consider ways of addressing this problem. In case of a co-variate study, the textbook suggestion for addressing omitted variable bias is to use instrumental variable analysis (Greene 2003). When we have an explanatory variable (X) that is supposed to have a causal effect on observed group differences in preferences (Y), the approach is to instrument variable X with a variable Z that affects variable X but has no direct effect on the dependent variable Y. Instrument Z would be a valid instrument,

meaning an instrument that can effectively address omitted variable bias, if it satisfies the following two criteria:

1. The instrument Z correlates with explanatory variable X , what is known as the *inclusion* restriction;
2. The instrument Z does not correlate with any other variables affecting the dependent variable Y , what is known as the *exclusion* restriction.²

In case of an identity-priming study, in contrast, the potential for omitted variable bias is much less recognized and a textbook solution absent. However, the ideas underlying instrumental variable analysis appear illuminating also for thinking about omitted variables bias and possible threats to clean identification in identity-priming studies. Identifying the causes of group differences in preferences—which is different from showing any causal effect of group membership on individuals' preferences—requires that we are able to identify exactly which group-level trait or traits are driving the observed treatment effect of the group identity prime. We think that a group identity prime is, in principle, able to provide such clean identification of causes. However, we also think that this requires that the prime only affects behavior of individuals belonging to different groups through the specific channel or trait proposed by the researcher. This requirement, in turn, suggests two criteria that need to be satisfied, akin to the criteria determining the validity of an instrument in instrumental variable analysis:

1. The identity prime needs to affect the salience of a specific trait of the group, for example, a specific cultural trait;
2. The identity prime may only affect the salience of this specific trait of the group and not one or more other traits of the group.

Emphasizing the link between these criteria and instrumental variable analysis, we can think of these two criteria as an *inclusion criterion for identity primes* (Criterion 1) and an *exclusion criterion for identity primes* (Criterion 2). More importantly, we think that, together, these two criteria help set the stage for

² In formal terms, we can think of the model $Y = \beta_0 + \beta_1 X + \varepsilon$ and state these two restrictions as $Cov(Z, X) \neq 0$ (inclusion) and $Cov(Z, \varepsilon) = 0$ (exclusion), where ε is the error term.

thinking systematically about how identity primes can substantiate causal inferences concerning the sources of group differences in preferences. As indicated, making group identity salient works like a black box and typically does not involve specific group-level traits that help explain observed differences in preferences between groups (Leung and Van de Vijver 2008; Schwarz and Strack 1981). However, as we discuss in more detail in the next section, considering the specific group-level trait(s) made salient by a particular group prime would be helpful in identifying as exactly as possible the source(s) of differences between groups.

To conclude this section let us note explicitly that the bar set by the proposed inclusion and exclusion criteria for group identity primes is rather high. Instrumental variable analysis already faces important challenges in establishing the validity of an instrument. The exclusion restriction in particular is difficult to satisfy, as any evidence showing that a given instrument satisfies this restriction is itself subject to omitted variable bias. For identity primes, however, the challenge would be higher still, as in this case also satisfaction of the inclusion criterion is not necessarily easily shown.

4. Group Identity Primes and Clean Identification in Practice

4.1. Distinguishing between alternative explanations for identity-priming effects

Our goal in proposing an inclusion and exclusion criterion for identity primes is not to invalidate prior work but to facilitate thinking about identification and the making of causal inferences about the sources of group differences in preferences. The essence of the potential for omitted variable bias in identity-priming studies is that room remains for multiple inferences concerning the factor that is driving the observed treatment effect of the identity prime. If the goal is rigorous measurement of group differences in preferences, a quasi-experimental research design suffices and there is no point in creating randomization using an identity prime (cf. Burnham and Kurzban 2005). Similarly, if the goal is not explanation but to show that group membership has a causal effect on individuals' preferences, a simple identity prime suffices and there is no need to think about whether or not the inclusion and exclusion restrictions for identity primes are satisfied. Oftentimes, however, the goal is to explain differences in

preferences between groups, in which case one needs to be able to identify specific causal factor and rule out other factors as driving the results.

To illustrate the issue let us first use a hypothetical example and then a real-world example. For the hypothetical example, consider Asian-American differences in preferences in a sample of undergraduate students at a university in the U.S. (cf. Benjamin et al. 2010). Aggregating experimentally elicited behaviors and comparing averages, we can determine the degree to which these two groups have different preferences concerning, for instance, risk. In addition, considering the treatment effect of a group-identity prime we can establish that the observed differences in preferences are, in fact, causally due to individuals' membership of either the Asian or the American group. What cannot be established merely on the basis of the observed effect of an Asian or American identity prime, however, is the specific source of the effect of being Asian or American on individuals' preferences. Culture, for instance, would seem a theoretically highly plausible explanation for observed Asian-American differences. Alternative (non-cultural) explanations are also possible, however, and cannot be rejected using the available evidence, i.e., on the basis of the observed effect of the identity prime. Consider, for instance, the possibility that the identity prime triggers individuals' awareness of the social status of their group relative to native-borns, which are the dominant group in terms of social status. In this case, we may find an effect on risk preferences but not (only) because of Asian culture but (also) because making their low social status salient affects people's willingness to take risks, as found by Mishra et al. (2015) and Payne et al. (2017), among others.

The second, real-world example helps show that the challenge to distinguish between alternative explanations for found priming effects is not just theoretical but both scientifically and societally relevant. The real-world example that seems most suitable for this purpose is Cohn et al.'s (2014) study of the effect of bankers' professional identity on dishonesty that we mentioned earlier. The reason is twofold, namely that Cohn et al.'s (2014) study speaks to a prominent societal debate—the business culture in the financial sector and its role in the global financial crisis (Financial Crisis Inquiry Commission 2011; Stiglitz et al. 2010; Van Hoorn 2015a)—and that this study has received much critical follow-up (Stöckl

2015; Van Hoorn 2015b; Vranka and Houdek 2015). The concern in the latter studies is that Cohn et al.'s (2014) main claim that the culture that exists in the banking sector promotes dishonest behavior is not necessarily substantiated by the finding that priming bankers' identity increases cheating in a coin tossing game. More in particular the concern is that this identity prime does not provide clean identification of an effect of business culture, while various alternative explanations of the observed priming effect appear at least equally plausible. Both Van Hoorn (2015b) and Vranka and Houdek (2015), for instance, propose that priming bankers' professionals identity activates a negative stereotype about the profession's honesty, which, in turn, is likely to have a negative effect on honesty in a laboratory game. As Cohn et al. (2013, p. 1) put it in an early version of the study published as Cohn et al. (2015): "a person's identity can change if society treats him or her as a criminal, leading to the adoption of behavioral propensities consistent with the criminal label." Meanwhile, it seems unlikely that the treatment effect of making inmates' criminal identity salient reported by Cohn et al. (2013, 2015) is a cultural phenomenon that is caused by the specific culture that exists among inmates.

4.2. Priming for explanation

The two examples above naturally raise the issue of what can be done to strengthen clean identification of the causes of group differences in preferences using identity primes. This issue is also the point of the inclusion and exclusion criterion for identity primes that we introduced in the previous section. Hence, the short response to this issue is that we cannot rely on observing a treatment effect of an identity prime alone and also need to assess the degree to which the particular prime used satisfies the inclusion and exclusion criterion for identity primes. The more plausible it is that the identity prime affects the salience of a specific trait of the group and only this trait, the stronger the evidence that this specific trait is a significant cause of observed group differences. Still, how to go about making this assessment in practice is an important open question. Hence, to complete the discussion of clean identification and explanation of group differences in preferences presented in this paper, below we consider some specific things that would help argue that the inclusion and exclusion criterion for identity primes are sufficiently satisfied.

Focusing on the satisfaction of the inclusion criterion for primes, we find that an important piece of evidence that would increase confidence that this criterion is satisfied is if the prime is not generic but seems narrower, targeting a specific feature of a group. A generic prime would make the whole of individuals' Asian identity salient, for instance, while a narrower prime might succeed in making only distinctively Asian collectivistic cultural tendencies salient (see Oyserman and Lee 2008 for a review of research using group primes that target specific dimensions of societies' culture). Considering the nature of the prime, generic vs. specific, researchers can be more or less confident about the channel through which the prime affects behavior and hence about the causes of the observed effect of the identity prime. More definitive evidence on the satisfaction of the inclusion criterion would involve additional data showing that the prime indeed makes particular group-level traits salient (in addition to affecting individuals' behavior or preferences). The study by Payne et al. (2017) is a nice example of how to go about this. They hypothesized that low social status increases risk taking because upward social comparison affects individuals' perceived needs and designed their research so that it involved experimental evidence not only of manipulated social status affecting risk taking but also of manipulated social status affecting perceived needs (measured using self-reports). In principle, any identity-priming study can collect additional data, notably self reports, that could help identify the specific group-level trait that is made salient by the prime used, thus strengthening the study's inferential power.

Focusing on the satisfaction of the exclusion criterion for primes, we find that increasing confidence that this criterion is satisfied is best served by collecting additional data. As is the case for the exclusion restriction in instrumental variable analysis, satisfaction of the exclusion criterion for group identity primes cannot be proven. Drawing on additional data it is possible, however, to show that a particular identity prime does not have an effect on certain group traits. When interested in the effect of Asian culture on risk preferences, for instance, it is possible to examine how the Asian identity prime affects subjects' subjective social status or perceived needs, among others. If the prime does not affect such measures, while the prime does affect risk preferences, we can safely conclude that the observed treatment effect of the prime does not occur because the prime increases subjects' awareness of their low

social status vis-à-vis native-borns. Overall, we find that there are several feasible and convincing ways in which identity primes can be relied upon to argue specific explanations of differences in preferences between pre-existing groups. Hence, even though there is plenty reason to take the potential for omitted variable bias in identity-priming studies seriously, there is no need for pessimism.

5. Discussion

Being able to distinguish between different explanations for a phenomenon matters because different explanations support different views on how the world works. Whether, say, male-female differences in preferences are mostly the result of social norms concerning gender roles or mostly a biological phenomenon will shape how one views, for instance, the male-female gap in wages or authority attainment. Moreover, depending on the source of male-female differences in preferences, government policies aimed at reducing gender disparities are more or less effective. Setting a quota for female board members may, for instance, contribute to changing gender norms but is unlikely to overturn biological differences between the sexes.

Still, an obvious concern with using an inclusion and exclusion restriction for group identity primes is that these criteria are too strict, raising the bar for clean identification in the context of experiments involving pre-existing groups too highly. Specifically, a most natural argument against the usefulness of considering two criteria that originated as a way of strengthening the inferential power of observational studies is that this would lead to a disregard of the many other important advantages offered by experimental research designs vis-à-vis observational research designs. We agree in the sense that we do not advocate holding experimental studies to higher standards, while giving a pass to studies facing even worse problems in identifying the causes of group differences in preferences. We disagree, however, in the sense that we do not think that having important advantages over other methods means we should not reflect on the use of identity primes for causal identification at all. Rather than invalidating identity-priming studies, the inclusion and exclusion restriction for group identity primes should facilitate systematic thinking about the conditions under which researchers can rely on priming to identify specific

causes of group differences in preferences. More generally, this paper is not meant as the final answer on how to go about the experimental study of group differences in preferences. Omitted variable bias appears an important issue in the experimental study of group differences between pre-existing groups and we therefore only intend this paper as providing a helpful starting point for future work contributing to the pushing of the causal inferences substantiated by behavioral experiments involving different human groups.

REFERENCES

- Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2008). Eliciting Risk and Time Preferences. *Econometrica*, 76, 583-618.
- Beck, A., Kerschbamer, R., Qiu, J., & Sutter, M. (2014). Car mechanics in the lab—Investigating the behavior of real experts on experimental markets for credence goods. *Journal of Economic Behavior & Organization*, 108, 166-173.
- Bellemare, C., Kröger, S., & Van Soest, A. (2008). Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities. *Econometrica*, 76, 815-839.
- Benjamin, D. J., Choi, J. J., & Strickland, A. J. (2010). Social Identity and Preferences. *American Economic Review*, 100, 1913-1928.
- Benjamin, D. J., Choi, J. J., & Fisher, G. (2016). Religious identity and economic behavior. *Review of Economics and Statistics*, 98, 617-637.
- Berry, J. W. (2002). *Cross-cultural psychology: Research and applications*. Cambridge: Cambridge University Press.
- Bond, M. H., & Van de Vijver, F. J. R. (2011). Making scientific sense of cultural differences in psychological outcomes: Unpackaging the magnum mystery. In: Matsumoto, D., & Van de Vijver, F. J. R., Eds. *Cross-cultural research methods in psychology*. Cambridge: Cambridge

- University Press: 75-100.
- Burnham, T. C., & Kurzban, R. (2005). On the limitations of quasi-experiments. *Behavioral and Brain Sciences*, 28, 818-819.
- Camerer, C., & Fehr, E. (2004). Measuring social norms and preferences using experimental games: A guide for social scientists. In: Henrich, J., et al., Eds. *Foundations of Human Sociality*. Oxford: Oxford University Press.
- Cason, T. N., Saijo, T., & Yamato, T. (2002). Voluntary participation and spite in public good provision experiments: an international comparison. *Experimental Economics*, 5, 133-153.
- Charness, G., & Villeval, M. C. (2009). Cooperation and Competition in Intergenerational Experiments in the Field and the Laboratory. *American Economic Review*, 99, 956-978.
- Cohn, A., Fehr, E., & Maréchal, M. A. (2014). Business culture and dishonesty in the banking industry. *Nature*, 516, 86-89.
- Cohn, A., Maréchal, M. A., & Noll, T. (2013). Bad boys: The effect of criminal identity on dishonesty. University of Zürich, Working Paper, (132).
- Cohn, A., Maréchal, M. A., & Noll, T. (2015). Bad boys: How criminal identity salience affects rule violation. *Review of Economic Studies*, 82, 1289-1308.
- Cook, T. D., & Campbell, D. T. (1979). *Quasi-experimentation: Design and analysis issues for field settings*. Chicago: Rand McNally.
- Croson, R., & Buchan, N. (1999). Gender and culture: International experimental evidence from trust games. *American Economic Review*, 89, 386-391.
- Croson, R., & Gächter, S. (2010). The science of experimental economics. *Journal of Economic Behavior & Organization*, 73, 122-131.
- Eckel, C. C., & Grossman, P. J. (2008). Differences in the Economic Decisions of Men and Women: Experimental Evidence. In: Plott, C. R., & Smith, V. L., Eds. *Handbook of Experimental Economics Results, Volume 1*. New York: Elsevier: 509-519.
- Fehr, E., & List, J. A. (2004). The hidden costs and returns of incentives—trust and trustworthiness

- among CEOs. *Journal of the European Economic Association*, 2, 743-771.
- Financial Crisis Inquiry Commission. (2011). *The financial crisis inquiry report: Final report of the national commission on the causes of the financial and economic crisis in the United States*. http://fcic-static.law.stanford.edu/cdn_media/fcic-reports/fcic_final_report_full.pdf.
- Gächter, S., Herrmann, B., & Thöni, C. (2010). Culture and Cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365, 2651-2661.
- Gächter, S., & Schulz, J. F. (2016). Intrinsic honesty and the prevalence of rule violations across societies. *Nature*, 531, 496-499.
- Gneezy, U., Leonard, K. L., & List, J. A. (2009). Gender Differences in Competition: Evidence from a Matrilineal and a Patriarchal Society. *Econometrica*, 77, 1637-1664.
- Gong, B., Yan, H., & Yang, C. L. (2015). Gender differences in the dictator experiment: evidence from the matrilineal Mosuo and the patriarchal Yi. *Experimental Economics*, 18, 302-313.
- Greene, W. H. (2003). *Econometric Analysis* (Fifth ed.). Upper Saddle River: Pearson Prentice-Hall.
- Guala, F. (2008). Paradigmatic Experiments: The Ultimatum Game from Testing to Measurement Device. *Philosophy of Science*, 75, 658-669.
- Haigh, M. S., & List, J. A. (2005). Do professional traders exhibit myopic loss aversion? An experimental analysis. *Journal of Finance*, 60, 523-534.
- Harbaugh, W. T., Krause, K., & Vesterlund, L. (2002). Risk Attitudes of Children and Adults: Choices over Small and Large Probability Gains and Losses. *Experimental Economics*, 5, 53-84.
- Henrich, J., et al. (2001). In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies. *American Economic Review*, 91, 73-78.
- Herrmann, B., Thoni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science*, 319, 1362-1367.
- Leung, K., & Van de Vijver, F. J. R. (2008). Strategies for strengthening causal inferences in cross cultural research: The consilience approach. *International Journal of Cross Cultural Management*, 8, 145-169.

- Marwell, G., & Ames, R. E. (1981). Economists free ride, does anyone else?: Experiments on the provision of public goods, IV. *Journal of Public Economics*, 15, 295-310.
- Matsumoto, D., & Yoo, S. H. (2006). Toward a new generation of cross-cultural research. *Perspectives on Psychological Science*, 1, 234-250.
- Mishra, S., Hing, L. S. S., & Lalumiere, M. L. (2015). Inequality and risk-taking. *Evolutionary Psychology*, 13(3), 1474704915596295.
- Oosterbeek, H., Sloof, R., & Van de Kuilen, G. (2004). Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7, 171-188.
- Oyserman, D., & Lee, S. W. (2008). Does culture influence what and how we think? Effects of priming individualism and collectivism. *Psychological Bulletin*, 134, 311-342.
- Payne, B. K., Brown-Iannuzzi, J. L., & Hannay, J. W. (2017). Economic inequality increases risk taking. *Proceedings of the National Academy of Sciences USA*, 114, 4643-4648.
- Plott, C. R., & Smith, V. L., Eds. (2008). *Handbook of Experimental Economics Results, Volume 1*. New York: Elsevier.
- Poortinga, Y. H. (2016). Integration of Basic Controversies in Cross-cultural Psychology. *Psychology and Developing Societies*, 28, 161-182.
- Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., & S. Zamir, S. 1991. Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study. *American Economic Review*, 81, 1068-1095.
- Schwarz, N., & Strack, F. (1981). Manipulating salience: Causal assessment in natural settings. *Personality and Social Psychology Bulletin*, 7, 554-558.
- Shadish, W. R., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. Boston / New York: Houghton Mifflin Company.
- Stiglitz, J. E., & Members of a UN Commission of Financial Experts. (2010). *The Stiglitz report: Reforming the international monetary and financial systems in the wake of the global crisis*. New York: New Press.

- Stöckl, T. (2015). Dishonest or professional behavior? Can we tell? A comment on: Cohn et al. 2014, Nature 516, 86-89, "Business culture and dishonesty in the banking industry". *Journal of Behavioral and Experimental Finance*, 8, 64-67.
- Van de Vijver, F. J. R., & Leung, K. (1997). *Methods and data analysis for cross-cultural research*. Thousand Oaks: Sage.
- Van Hoorn, A. (2012). Cross-cultural experiments are more useful when explanans and explanandum are separated. *Proceedings of the National Academy of Sciences USA*, 109, E1329-E1329.
- Van Hoorn, A. (2015a). The global financial crisis and the values of professionals in finance: An empirical analysis. *Journal of Business Ethics*, 130, 253–269.
- Van Hoorn, A. (2015b). Organizational culture in the financial sector: Evidence from a cross-industry analysis of employee personal values and career success. *Journal of Business Ethics*. doi:10.1007/s10551-015-2932-6.
- Visser, M. S., & Roelofs, M. R. (2011). Heterogeneous preferences for altruism: gender and personality, social status, giving and taking. *Experimental Economics*, 14, 490-506.
- Von Gaudecker, H. M., Van Soest, A., & Wengström, E. (2011). Heterogeneity in risky choice behavior in a broad population. *American Economic Review*, 101, 664-694.
- Vranka, M. A., & Houdek, P. (2015). Many faces of bankers' identity: how (not) to study dishonesty. *Frontiers in Psychology*, 6, 302.