Identity and Choice Under Risk

Francesco D'Acunto *

January 2020

Abstract

Traditional gender norms have been blurring over the last years. Could this evolution have implications for economic decision-making? Identity theory posits that men who commit more to traditional male identity norms should take on more risk through a subjective-beliefs channel when their identity is salient, either because it is primed or threatened. I test these predictions using large-scale artefactual field experiments. Men whose identity is primed or threatened invest more in risky opportunities than control men and women. They become overconfident even in pure games of chance with no scope for skill, which is consistent with the subjective-beliefs channel identity theory postulates. The effects are stronger for men who commit more to male identity—older men and men in the Southern US. The recent blurring of traditional gender norms might thus imply a drop in aggregate risk taking over time, as men might reduce their willingness to take risks in disparate contexts.

JEL: C91, C93, D03, D14, G02, J16, Z13

Keywords: Cultural Economics, Subjective Beliefs, Overconfidence, Norms, Stereotypes, Financial Decision-Making, Heterogeneous Agents.

*Carroll School of Management, Boston College (dacuntof@bc.edu). The experimental procedures in this paper have been authorized by the IRBs at the University of California at Berkeley (CPHS Approval Protocol number 2013-11-5805) and the University of Maryland (UMCP IRB Protocol number 842799-1). For their invaluable help and guidance on this project, I thank Stefano DellaVigna, Ulrike Malmendier, and Gustavo Manso. For very helpful comments and discussions, I thank Nick Barberis, Dan Benjamin, James Choi, Xavier Gabaix, Luigi Guiso, George Korniotis, Samuli Knüpfer, Rachel Kranton, Pete Kyle, Ross Levine, Don Moore, Adair Morse, Muriel Niederle, Terry Odean, Christine Parlour, Matthew Rabin, Yuval Rottenstreich, Kelly Shue, Stephan Siegel, Michael Weber, and seminar participants at the 2015 Western Finance Association, the 2015 Stanford Institute for Theoretical Economics (Experimental Economics), the 2014 Miami Behavioral Finance Conference, the 2014 NBER Behavioral Economics Fall meeting, the 2014 Whitebox Advisors Graduate Conference, and UC Berkeley. Financial support from the White Foundation is gratefully acknowledged. All errors are my own.

Over the last few decades, the norms of behavior attached to male and female identities have been converging substantially (e.g., see Ciabattari (2001); Glaeser and Ma (2013); Goldin (2014)). This phenomenon has reached unprecedented levels to the extent that the notion of a continuum of gender norms as opposed to binary genders has diffused swiftly (Richards et al. (2016)): Figure 1 shows that US men in younger cohorts see their gender as more fluid and more faceted in the masculine/feminine scale relative to US men in older cohorts (Moore (2016)).

Figure 1: **Blurring Gender Norms across Male Cohorts**

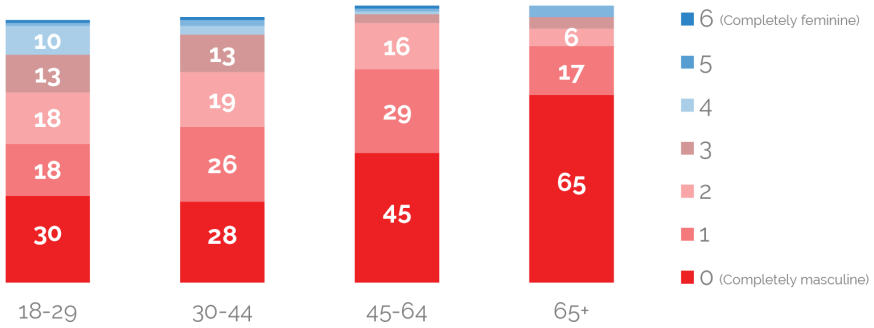


Figure 1 replicates Figure 1 in Moore (2016), which reports the answers to a large-scale YouGov.com poll among US men of various ages regarding the extent to which men see themselves as masculine, feminine, or intermediate gradations between the two binary genders.

Could this secular convergence in the perception of gender norms have implications for economic decision-making? Research in social psychology suggests so. Starting from Mead (1934), identity theory posits that individuals commit to several role identities, each prescribing norms of behavior that are common knowledge in society.[1] Individuals conform their choices to such norms to verify their identity. Verifying a role identity "means [...] enacting behavior consistent with the role" (Burke (1980b)). Choices that are incongruent with one's identity make individuals worse off, and hence individuals tend conform their choices to their role identities that are salient when making decisions.

Because traditional male identity norms prescribe risky and aggressive behaviors (see Schwartz and Tammy (2005) and Table 1), men should take on more risk when their identity is salient. Moreover, identity verification should affect men more than women, because "manhood, in contrast to womenhood, is seen as a precarious state requiring continual social proof and validation" (Vandello et al. (2008)). And, these effects should

[1]See Stryker (1980), Burke (1980a), Eagly (1987), and Carli (1989), among others

be stronger for men that commit more to traditional identity norms, whereas they should decline and ultimately fade for men whose gender identity norms blur and blend in a continuum with female identity, as seems to be the case for the most recent cohorts.

In this paper, I aim to test empirically for these and other unique predictions of identity theory for choice under risk, to dig into the economic channels through which identity might affect risky choices, and to assess the extent to which burring gender identity norms might modify well-known regularities in choice under risk across genders.

The starting point is testing whether the basic predictions of identity theory are confirmed in the field. Do men whose identity is salient react by taking on more risk? According to identity theory, the *same* reaction should arise when male identity is either primed or threatened, because in both cases men are better off when they can verify/reaffirm their own identity and avoid incongruence of their action with identity norms.[2]

That men should take on more risk when their identity is threatened is a unique prediction of identity theory, which allows for a falsification test of this theory. A direct effect of identity salience on risk taking might seem a trivial prediction: Once men are reminded about situations in which they behaved "manly," that is, in risky and aggressive ways, they might take on more risk for many reasons. Such triviality, though, does not arise for identity threats, i.e. when men think about situations in which they behaved in line with female identity norms. Conformity to externally induced norms of behavior or experimenter demand effects—subjects' attempt to guess the experimenter's aims and to behave in line with such aims—would if anything predict that men should take on *less* risk when primed with female identity, because female identity is associated with non-risky norms of behavior.

Moreover, based on identity theory women should react to either male of female identity primes less than men, because incongruence with one's identity is more detrimental to men than to women (Vandello et al. (2008)). Women's differential reaction is instead inconsistent with conformity or demand effects. If subjects were merely taking on more risk because they were primed with risk taking, men and women should both take on more risk when primed with male identity, which is not what identity theory hypotesizes.

To obtain exogenous variation in the salience of gender norms, I test for these baseline

[2]This prediction echoes the male overcompensation hypothesis in social psychology (e.g., see Maas et al. (2003); Willer et al. (2013).

predictions in a set of artefactual field experiments. I randomly prime or threaten subjects' gender identity (Steele (1997); Shih et al. (1999)).[3] The prime consists of asking men and women to read a (true) internet blog, which describes the typical social stereotypes society attaches to men or women's behavior in line with those in Table 1. After reading, subjects write a short essay (5-10 sentences) where they describe in detail a situation in which they behaved in line with the norms discussed in the text. For the identity threats, the experimental conditions are reversed—men read the female-identity text and recall and describe a situation in which they behaved in line with feminine norms, whereas women write the essay after reading the male-identity text.

I first elicit subjects' risk attitudes in a within-subject design both before and after the experimental manipulation using incentivized lottery choices à la Holt and Laury (2002). Consistent with identity theory, I find that both groups of men whose identity is primed or threatened choose riskier lotteries after the treatment relative to before, which is not the case for control men or for women.[4]

After having established the baseline effects of identity norms of choice under risk, I move on to assess the channels of transmission. A crucial prediction of identity theory is that the effect should involve a beliefs channel—identity salience should produce motivated beliefs (Stets and Burke (2000); Benabou (2015)). Because identity verification makes men better off, men should increase their subjective probability of experiencing good states if they accepted a risky bet to induce higher risk-taking than might be optimal based on objective probabilities. In the identity-theory language of Stets and Cast (2007), "individuals manipulate beliefs as personal resources, including one's worth and efficacy. Positive and efficacious self-feelings help individuals be more persistent in the face of difficulties in verifying the self." A unique consequence is that the effect of male identity on beliefs should arise even in pure games of chance with no role for skill and for which a neoclassical agent would hold subjective beliefs identical to objective probabilities, when known.

[3]Economic research has employed priming techniques to address questions for which detecting exogenous variation in the field is hard. For instance, see Benjamin, Choi, and Fisher (2013); Cohn, Engelmann, Fehr, and Marechal (2015); Coffman (2014)). Salience has been shown to affect choice under risk theoretically and empirically (e.g., Bordalo et al. (2012); Gargano and Rossi (2018); Frydman and Mormann (2018); Frydman and Wang (2018)).

[4]Identity theory posits that women obtain less utility than men from verifying their identity, and hence should react less than men to identity salience, but not necessarily that women should not react at all to the primes. Non-reaction, which is what I find, is consistent with identity theory, but a lower yet detectable reaction by women would have also been consistent with the theory.

To test for this unique prediction on the emergence of motivated beliefs, I design a task that elicits beliefs about risky outcomes without uncertainty abstracting from risk preferences. Subjects report what they think would be their probability of winning in an incentivized random-generating-process game whose known probability of winning is 50%. A standard agent would respond that the probability is 50%, because the game contains no uncertainty, the probability of winning is known, and the random-generating-process allows for no role of individual abilities and skills to influence the outcome. Men whose identity is primed or threatened, instead, think that they would be more likely to win than 50% of the times. At the same time, the same primed men do not hold distorted beliefs about the probability that their peers would win in the same situation, which is prima facie evidence that primed men understand the objective probabilities that are stated to them but believe they are better than others even in pure games of chance. Control men and women do not believe that they (or their peers) are more likely to win than 50% of the time.

The second unique prediction of identity theory in terms of channels is that the norms attached to male identity, and not the gender context in itself, should drive the effects. This point implies that *any* other social identity that prescribes risky norms of behavior similar to male identity should affect men's risky choices similarly, even if gender is not mentioned. To test this prediction, I make salient a powerful role identity, which, consistent with earlier research, I show prescribes norms of behavior that are almost indistinguishable from those male identity prescribes.[5] And, indeed, men primed with power identity react in a qualitative and quantitative similar way to men primed with male identity. Similar to gender identity, power identity has no effect on women.

To verify that the framing of the tasks does not drive the results, and to assess whether the baseline effects on risky lottery choices might also arise in more complex economic decisions, I study two applications to financial decision-making regularities broadly documented in the literature. First, male identity norms should increase men's risk-taking not only when facing positive expected value (EV) opportunities—for instance, investing in an index fund—but also when facing negative EV opportunities—for instance, excessive trading of individual stocks, because identity verification makes men better off and the gains from identity verification might offset the consumption disutility

[5]Research in social psychology has used power-identity primes to induce overconfidence in the form of illusion of control in human subjects (e.g., see Anderson and Galinsky (2006) and Fast et al. (2012)).

from investment losses. Even risk-averse agents, who under expected utility theory should shy away from negative EV investments, might invest in such opportunities because of the benefits of identity verification.[6] I find evidence consistent with this prediction in an experiment in which subjects face opportunities framed as financial investments whose EV varies around zero.

The second application I consider is delegated choice under risk, which is typical of managerial investment decisions or financial advisers' management of households' wealth. According to agency theory, agents might take on excessive risk when managing a principal's resources than when managing their own resources if incentives are misaligned. Instead, beliefs-based arguments such as the overconfidence hypothesis propose that delegated agents might take on excessive risk because they mistakenly believe their excessively risky choices are optimal for their principal.[7]

Identity theory suggests an additional channel: Delegated agents might take on excessive risk to enjoy the benefits of identity verification. Under this explanation, identity should affect men's choices similarly irrespective of whether choices are direct or delegated. Accordingly, male identity has similar effects on the risk-taking behavior of men who make decisions framed as direct or delegated investments.[8]

After having tested the main predictions of identity theory for choice under risk empirically, I can finally tackle the opening question: Could converging male and female identity norms have implications for men's risk taking? Male identity should affect men who are highly committed to traditional male identity norms more than other men, because the utility benefits of identity verification are higher for identities to which individuals strongly commit (Burke (1980b)).[9]

Because asking subjects directly about their commitment to traditional gender norms

[6]For a discussion on the effects of cultural norms on beliefs about investing in financial markets see D'Acunto, Prokopczuk, and Weber (2018)

[7]For instance, see Malmendier and Tate (2005), Malmendier and Tate (2008), Hirshleifer, Low, and Teoh (2012), and Ben David, Graham, and Harvey (2013).

[8]Unfortunately, the experimental setting does not allow me to create actual delegated choices and hence the direct and delegated investment choices vary in the framing of the situation. Earlier research in social psychology, starting with Tversky and Kahneman (1981), shows that framing has a substantial impact on choice in otherwise similar settings.

[9]But see also Benjamin et al. (2010), who discuss an ambiguous prediction of identity for the relationship between commitment and reaction to identity salience.

would have primed all subjects with such norms,[10] and I could not run the experiments over several decades while gender norms were blurring, I tackle this question using cross-sectional variation outside the experiments in subjects' commitment to traditional male identity norms.[11]

First, I consider men who belong to different cohorts. Not only have gender-identity norms converged over the last decades, but the dichotomy between masculine and feminine stereotypes also tends to become starker with age (Biernat (1991)). Men in older cohorts should thus commit more to traditional male identity norms than men in younger cohorts, because of both a cohort and an age effect, both of which effects are consistent with the patterns in Figure 1. And, indeed, I find that men in older cohorts react to male-identity salience by taking substantially more risk, but the effects fade monotonically for younger cohorts to the extent that men in the youngest cohort do not react at all. This result obtains in all the choice framings described above, including the lottery choices, the motivated-beliefs channel, and financial investment in negative EV opportunities and by delegation.

Second, I consider identity commitment across space by comparing the effects of identity salience on men who live in the Southern US and other men. Men in the US South tend to have more conservative views about gender norms than others, based on survey evidence such as the answers to questions in the General Social Survey (e.g., Rice and Coates (1995); Powers et al. (2003); Davis and Greenstein (2009)). Consistently, the effects of identity on men's choice under risk are substantial for individuals in the South but smaller in other US regions, although I do not detect statistical significance across all sample splits.

Overall, the results on the heterogeneity by commitment to gender norms are consistent with the possibility that fading traditional gender norms might reduce men's willingness to take risks in economic contexts. Note that the interpretation of the evidence in this paper has caveats. By construction, the analysis isolates reduced-form effects. The experiments change the salience of identity cues to gauge their causal effect on

[10] In fact, indirect elicitation methods of subjects' gender stereotypes exist, such as the Implicit Association Test (IAT), which has been used successfully in research on gender stereotyping (e.g., see Carlana (2019)).

[11] In fact, these tests were not part of the initial set of hypotheses, and hence I did not stratify the subjects I recruited in order to maximize the power of these heterogeneity tests. When discussing the results of these tests, I emphasize how the lack of ex-ante stratification reduces statistical power.

decision-making, but the perception level of cues affects decisions in the field (e.g., see Ravina (2008); Barone et al. (2015); and Ravina (2017)). Caveats also include the fact that the choices contain no uncertainty and the within-subject design combined with priming raises concerns about demand effects. In Section 6, I discuss a direct test and indirect arguments suggesting demand effects do not drive the results in this paper.

Stressing what this paper does not show is also important. The paper does not show that gender identity is more important than biology, neurology, or genetics as a determinant of choice under risk. In fact, any of these dimensions could mediate the effects of identity on risk attitudes and beliefs. Moreover, the results are inconsistent with the idea that manipulating gender identity could make the risk attitudes of men and women identical, because men take on more risk once primed in any direction, whereas women do not react.

An intriguing implication of identity theory is that acting on biased beliefs might make agents better off than acting on objective beliefs. Choices inconsistent with expected utility theory are not necessarily suboptimal in this context.[12] For instance, for the case of excessive trading of individual stocks, the utility benefits stemming from identity verification might overcome the expected losses of investing in typically negative EV opportunities. Nudges that guide agents to the optimal choices under expected utility theory are not necessarily beneficial in this identity theory of choice under risk. Providing a framework to compare the benefits of consumption utility with those arising from psychological sources to allow welfare statements is a challenge for future research (e.g., see Guthrie and Sokolowsky (2013)).

1 The Controlled Environment

I run the experiments on an online platform, Amazon Mechanical Turk (mTurk), which researchers in economics are using with increased frequency.[13] On mTurk, Requesters post tasks, and a large pool of Workers can accept to perform them. Requesters are often private companies, and Workers are registered users. Workers provide their fiscal address and social security number for tax purposes. Tasks are short, and the average pay is low

[12] See also Exley and Kessler (2018), who show that well-known behavioral patterns attributed to behavioral biases disappear once they eliminate self-serving motives. D'Acunto (2018) and D'Acunto et al. (2019c) find motivated beliefs can help understand limited risk-taking across demographic groups.

[13] For instance, see Kuziemko et al. (2015); Bordalo et al. (2018); Landier, Ma, and Thesmar (2017); DellaVigna and Pope (2018); Ferman et al. (2018); Lian, Ma, and Wang (2018); D'Acunto et al. (2019a); and D'Acunto et al. (2019b).

($1.39 per hour). Workers based in the US, whom I exclusively recruit in this paper, access mTurk mainly to spend their spare time constructively (Paolacci et al. (2010)).

The quality of answers and task performance of Workers on mTurk appears to be similar, if not better, to those of human-subjects laboratories despite the lower pay (Casler et al. (2013)).[14] This fact should not appear surprising, because the incentives on an internet platform like mTurk are different from those in a human-subject laboratory in at least two respects. First, the apparently low monetary values are in line with the average expected earnings of subjects on the platform. Second, and crucially, mTurk provides recognition for the best-performing Workers ("Master Workers"), based on Requesters' reviews. This recognition is an important incentive scheme on a social platform like mTurk and has no parallel in human-subjects laboratories.

Nonetheless, to reduce concerns about Workers' lack of effort: (i) I restrict the subject pool to Workers with at least 95% positive rates on all the tasks accepted in the past, (ii) I track the time they take to complete each task, (iii) I read all the short essays subjects have to produce as part of the experimental procedures described below to verify that subjects write coherent statements, and (iv) I add implausible options to the lotteries and I verify they are not picked.[15]

A Advantages over Human-subject laboratory for a Culture Priming Study

For a study on cultural-based priming and for within-subject designs, mTurk has a set of specific advantages over human-subject laboratories (see also Horton et al. (2011)):

- Subjects come from the whole US. Figure 1 plots their IP addresses. This feature is crucial to test for the variation in the effects of identity on choice under risk across geographic areas in which commitment to male-identity norms varies. This feature also guarantees a representative sample in terms of the meaning of role identities. If subjects came from the same college town, they would live in a peculiar social environment that is likely dissimilar from the one average Americans face. Also, students may sort into communities whose values they share.
- Subjects' demographics are well varied (see Figure B.1 of Online Appendix B), whereas a sample of college students would only include individuals between 18 and

[14] Mason and Suri (2012) describe mTurk and Workers. Berinsky et al. (2012) replicate laboratory results in political science on mTurk.

[15] For instance, I add a choice between $0 for sure and a lottery that pays $0 or a positive amount.

22 years of age with some college education. Demographic variation allows assessing the differences in the effects of identity on choice under risk across cohorts, which would not be possible had the subjects been college students.

- As an online, double-blind platform, mTurk allows the running of artefactual field experiments in which recruiting is simple (List (2011)). Subjects perform tasks in their environment and are not aware they are part of an experiment. This procedure helps me address the concerns of demand effects that within-subject designs raise.
- Replicability of results is easy. Any Requester accesses the same subject pool. Easy replicability allows for a transparent comparison of results across studies.[16]
- Workers' anonymity makes the priming procedure most effective, because subjects can describe unconventional experiences they might have faced or imagined as part of the priming task, with no fear of being identified by the experimenters.

2 Identity and Risky Choices

The norms of behavior attached to male identity include attitudes related to choice under risk, such as aggressiveness, risk taking, strength, confidence, assertiveness, and dominance. I document this fact in Table 1, where the left column reports the most cited typical behaviors a set of 200 online respondents—unrelated to the subjects in the experiments I will describe later—associate with male identity. Based on identity theory, the norms of behavior attached to each identity are common knowledge in society, and hence asking a random sample from a pool with varied demographic characteristics should allow gauging this common knowledge.

The two baseline predictions of identity theory for choice under risk are that higher salience of male identity should induce riskier choices in men, and that this effect should be true for both priming and threatening male identity. They derive from the notion of identity verification (Burke (1980b)). According to this notion, individuals conform their behavior to the norms they associate with the identity that is most salient among those to which they commit. In so doing, individuals confirm their identity, which avoids the drop in utility they would face if they experienced any incongruence between their choices and their identity. Individuals who verify their identity have higher utility, ceteris paribus, than individuals who do not verify their identity. Higher utility after identity

[16]On September 26, 2012, Daniel Kahnemann proposed a protocol for improving the credibility of research on priming, which is easy to implement on mTurk.

verification or lower utility when facing incongruence to one's identity can drive this effect.

For the case of men, identity verification requires conformity to masculine normative behaviors irrespective of how and why male identity is salient. As such, men should react in the same way—acting in line with masculine normative behaviors—irrespective of whether male identity is salient because it is primed or threatened.[17] If men aligned with feminine norms of behavior after female identity was made salient to them (identity threat) they would falsify identity theory, because they would not confirm their own identity. Instead, such behavior could be consistent with conformity or demand effects.

Assessing these basic predictions is important because previous research in different fields found conflicting results on the effect of gender identity on risk attitudes. In economics, Benjamin et al. (2010) found no effect of gender salience on risk preferences in an incentive-compatible, between-subjects experimental design. In social psychology, Meier-Pesti and Penz (2008) found male gender salience increases the willingness to take risks in both men and women, whereas female gender salience decreases risk taking in men. This latter result contrasts with the findings of the overcompensation literature (Willer et al. (2013)) and the predictions of identity theory, but was obtained in a non-incentive-compatible setting and using a subject pool in which men and women reported the same masculine attributes before the experiments.

To reconcile these conflicting results, I propose a within-subject experimental design that controls for the heterogeneity of preferences before the experimental manipulations. This design allows testing for the effect of identity on risk attitudes more precisely than between-subjects designs, in which the cross-sectional variation in subjects' preferences when they access the experiment reduces statistical power drastically. Even in my experiments, if I only use the between-subjects variation in risky choices after the manipulations without conditioning for preferences before the experiments I estimate effects of similar size but often not statistically different from zero, which is consistent with the non-result of Benjamin et al. (2010).

Moreover, my priming procedure is purposefully more aggressive than in Benjamin et al. (2010): Subjects read a text that includes stereotypical male or female behaviors, and they recall and describe in detail a situation when they behaved in line with such

[17]This prediction recalls the notion of *overcompensation* studied in social psychology (see Willer et al. (2013)).

stereotypes.[18] This procedure makes salient a specific set of norms of behavior attached to genders, as opposed to any other channels through which thinking about gender might affect subjects' choices and beliefs. In Section 6, I discuss why my experimenter demand effects are barely consistent with my results.

Experiment 1: Design. I recruited 340 subjects on mTurk in April 2012 (first session) and September 2012 (second session). The subject pool was restricted to users with a US tax identification number (SSN or ITIN), and with more than 95% of lifetime tasks approved in the past. The recruiting ad description proposed a creative-writing task for which subjects would earn \$0.50, plus a chance to earn a bonus by picking lotteries. Subjects knew the bonuses were computed dividing by 100 the nominal amounts they earned in each of the lottery choice tasks, and that bonuses could range anywhere between \$0 and \$0.42. The average time subjects took to complete the tasks was 15 minutes and 2 seconds, which implies a hourly salary between \$2.00 and \$3.68, depending on the lottery choices. The payments in Experiment 1 as well as all other experiments in this paper were designed to be in line with the average hourly pay of Workers on mTurk, to avoid any undue influence on Workers when choosing between the current study and other tasks among those available on mTurk (Paolacci et al. (2010)). Subjects had to complete the tasks in full before proceeding across screens, but they could leave the experiment at any time. No one left the experiment before completing it.

The nature and objectives of the study were not disclosed, to ensure (i) subjects were not induced into demand effects, (ii) subjects would not search for previous research on identity priming, and (iii) the results could be replicated in subsequent sessions and by other researchers.

Table 2 describes the experimental designs and samples for all the experiments in the paper. The design of Experiment 1 was a 2 (male, female) X 3 (control, male prime, female prime) factorial design. As is customary in studies that use Holt and Laury (2002) lotteries to elicit risk aversion, I excluded 20 subjects (5.8% of the full sample) who switched across the certainty equivalents and the lottery more than once, and hence did not understand the elicitation task. The final sample thus consisted of 320 subjects. I assigned subjects in proportions of 2:2:1 to the control, male-prime, and female-prime conditions.[19] The aim was to maximize the power for the male-identity-salience test,

[18] Online Appendix A provides the experimental materials.

[19] About 60% of US mTurk workers are women (Mason and Suri (2012)).

because previous research had not tested my deception-free identity-threat test. I did not use the threat in the overcompensation literature (Maas et al. (2003); Willer et al. (2013)) to avoid deception.[20]

Experiment 1: Procedure. In the first stage, subjects answered four background questions including their country (which I used together with the recorded IP addresses to verify consistency with the restriction of the analysis to US-based Workers), gender, age group (18-22, 23-35, 36-45, 46-60, 60+), and education (high school or lower, some college, college degree, postgraduate degree). Crucially, all subjects across all experimental conditions, including control subjects, answered the same questionnaire, which included stating one's gender. Therefore, any potential effect of the gender question is absorbed by the behavior of the control group in this design.

Then subjects faced two screens of lottery choices (Holt and Laury (2002)). I report the example of one screen of lottery choices in Online Appendix A. Each choice included a degenerate lottery paying a positive outcome for sure (certainty equivalent), and a lottery paying a positive outcome with probability 1/2, and zero otherwise.[21] In the second session, subjects faced three screens of lottery choices to ensure the results did not vary with effort, which was not the case. All screens differed in the certainty equivalents and the positive outcome of the degenerate lotteries as well as in the number of overall choices to be made (between 15 and 20) to avoid that subjects would merely pick always the same switching points throughout screens without considering the options carefully. All choices were incentive compatible. Subjects knew their bonus would be calculated by picking one screen and one line at random, running their choice (certainty equivalent or lottery), and dividing the final amount by 100.

The second stage consisted of the experimental treatments. Subjects first read a short text. The reader can refer to the full texts for each condition in Online Appendix A. The texts were taken from online blogs, so as to be similar to the gender-identity manipulations to which the subject pool, which consisted of internet users, may have been regularly exposed in daily life. This test did not aim to give subjects a scientific

[20] They gave random feedback on a gender-identity survey that suggested subjects are masculine or feminine, which implies an act of deception on the part of the experimenter.

[21] Andreoni and Sprenger (2012) discuss the problems that arise when jointly estimating the discounting and concavity parameters of non-linear utility functions using discontinuous experimental budgets. In this paper, I do not aim to estimate any utility parameters. The difference-in-differences design within subjects washes away any potential systematic effect due to the elicitation method.

definition of identity norms—which is not what the priming procedure is about—but to activate their *own* interpretation of identity stereotypes once they faced them explicitly in a context not dissimilar from their own environment. Subjects in the male- and female-identity-prime conditions read texts including several stereotypes about the prototypical behavior of a masculine and feminine person. Control subjects read a text on ayurveda principles for a healthy lifestyle. One might be concerned that the control condition by itself might have somehow affected subjects' willingness to take risks. Fortunately, this concern can be tested directly in a within-subjects design. As discussed below, control subjects made very similar choices before and after the manipulation, which rules out an effect of the control text on risk taking.

Consistent with Table 1, the male-identity text includes norms like "risk taking". One might worry that exposing men explicitly to a prime that discusses risk taking might mechanically make them take on more risk. But even here the predictions on identity threat are helpful. If the concern was relevant, men would not take on more risk after identity threats, in which they are not exposed to a norm of risk taking. Moreover, women would also take on more risk when reading the male identity text (identity threat for women), because they would be exposed explicitly to a norm of risk-taking. None of these two predictions is consistent with identity theory and the experiments I design allow for a direct test of these predictions, both of which I reject.

The male-identity text also includes "confidence," which is a masculine norm. At the same time, the text states explicitly that "The masculine side also includes the ability to [...] accurately weigh probabilities so that you know the most likely outcome to expect in situations you come across." If anything, the text induces to an accurate assessment of objective probabilities as an important masculine norm of behavior, which is the opposite of the subjective-beliefs channel identity theory postulates.

After reading the text, subjects had to recall and describe an experience in line with the text. Subjects were asked to describe the situation and their feelings in detail in a short essay of 5 to 10 sentences. Note that whether the experiences subjects described were authentic or not is not relevant. Even if subjects described a hypothetical situation, they would think about identity roles and their gender's identity norms would become salient to them, thus triggering the quest for identity verification.

The third stage consisted of two additional screens of lottery choices. The lottery choices

in the third stage differed from the first stage in terms of certainty equivalents and positive payoffs of the degenerate lotteries to ensure subjects did not merely repeat their first-stage choices. Incentive compatibility was as in the first stage.

Manipulation Check. Across all experiments, I used the content of the essays the subjects produced for the manipulation checks. I verified that each essay reported words related to the identity stereotypes enlisted in Table 1 in the corresponding experimental condition.[22]

Figure B.2 of Online Appendix B describes this manipulation check for Experiment 1. The (unreported) manipulation checks for the subsequent experiments are similar. In Panel (a) of Figure B.2, I show the average number of times subjects wrote a masculine norm listed in Table 1, across experimental conditions. Subjects in the male-identity-prime condition wrote on average 2.41 male identity norms. By contrast, subjects in the control or female-identity-prime condition did not write male identity norms. In Panel (b) of Figure B.2, I report the average number of times subjects wrote a feminine norm from Table 1 in their essays, across experimental conditions. Subjects in the female-identity-prime condition wrote on average 2.16 female identity norms, whereas subjects in other conditions did not. If subjects merely recalled some words from the text and used them in their essay, the manipulation would have worked, because its aim was to have subjects think about male- or female-identity norms of behavior.

Measuring Risk Tolerance and Non-parametric Analysis. To compute a non-parametric measure of risk tolerance at the subject level, I subtract the times the subject chose the certainty equivalent over the lottery from the times a risk-neutral agent would make this decision. I report a graphical depiction of this procedure in Figure B.3 of Online Appendix B. I average the difference across the choices of the first stage to obtain a pre-treatment measure of risk tolerance:

$$Risk\ Tolerance_{pre,i} = \frac{1}{L} \times \sum_{l} (RN\ choices_{pre,l} - choices_{pre,l,i}),$$

where $l \in (1, L)$ are the screens of lottery choices the subject faced, $RN\ choices_{pre,l}$ are the times a risk-neutral agent would choose the certainty equivalent over the lottery in screen l, and $choices_{pre,l,i}$ are the times agent i chose the certainty equivalent over the

[22] Ad-hoc manipulation-check tasks would have made subjects aware they were part of an experiment, and hence would have worsened the risk of demand effects.

lottery in screen l. $Risk\ Tolerance_{pre,i}$ is positive when the subject chose the lottery more often than a risk-neutral agent. I compute the analogous measure for choices made after the experimental treatments:

$$Risk\ Tolerance_{post,i} = \frac{1}{L} \times \sum_{l} (RN\ choices_{post,l} - choices_{post,l,i}).$$

The within-subjects change in risk tolerance after the treatment relative to before the treatment is $\Delta Risk\ Tolerance_i = Risk\ Tolerance_{post,i} - Risk\ Tolerance_{pre,i}$.

Figure 3 plots the estimated densities and the cumulative distributions of $Risk\ Tolerance_{pre}$ and of $\Delta Risk\ Tolerance$ for the subsample of men, who according to identity theory should respond to the experimental manipulations. Panel A of Figure 3 plots the estimated distributions of the level of risk tolerance of men across treatment conditions before the manipulations. The vast majority of subjects were risk averse. The mean of all distributions and the largest part of their mass lie in the negative domain. Reassuringly, the distributions for men that were exposed to the control and to the male-prime conditions (green, dot-dash and blue, solid lines) are similar in terms of mean and standard deviation. For men in the female-prime group, the standard deviation is slightly lower, but I cannot reject the null that all three distributions are the same statistically.

Panel B of Figure 3 depicts the main result of Experiment 1. The plots report the density and cumulative distribution of $\Delta Risk\ Tolerance$ for the male subsample across experimental conditions. The risk tolerance of control subjects (green, dot-dash line) does not change, on average, after exposure to the control condition. This evidence suggests the control condition does not affect the results, and the acts of recalling and writing do not cause any change in choices by themselves.

If the male prime and the female prime were both increasing men's risk tolerance as predicted by identity theory, we would expect the distributions of $\Delta Risk\ Tolerance$ for the treatment conditions to lie to the right of the distribution for controls. And, indeed, the change in the risk tolerance of the male-prime group (blue, solid line) is on average positive (0.5). The same is true for the male-identity group (dashed, red line; 0.45).

I report the analogous graphical evidence for the female subsample–which according to identity theory should not necessarily react, given that the utility stemming from identity

verification is lower for women than for men—in Figure B.4 of Online Appendix B.

OLS Results. To assess the statistical significance of the effects, I first estimate the following difference-in-differences specification by OLS on the subsample of men:

$$\Delta Risk\ Tolerance_{iae} = \alpha + \gamma_1 \times Male\ Prime_{iae} + \gamma_2 \times Male\ Threat_{iae} + \eta_a + \eta_e + \epsilon_{iae}, \quad (1)$$

where $\Delta Risk\ Tolerance_{iae}$ is the within-subject change in risk tolerance for subject i, in age group a, and education group e; $Male\ Prime$ and $Male\ Threat$ are dummies that equal 1 if the man was exposed to the male-identity (salience) or female-identity (threat) prime; η_a and η_e capture subjects' age and education groups. In terms of statistical inference, I report Huber-White standard errors, which are equivalent to standard errors clustered at the subject level given that equation (1) only includes one observation per subject.

The identification strategy is a difference-in-differences design: I compute subjects' change in risky choices before and after the manipulations, and across subjects in the treatment and control groups.[23] Column (1) of Panel A of Table 3 shows that men primed with male identity chose lotteries over certainty equivalents 0.51 times more often (s.e. 0.25) than controls after the prime. The result is robust to averaging out age- and education-group effects in column (2) of Panel A of Table 3. Men primed with female identity chose lotteries over certainty equivalents 0.43 times more often (s.e.0.21) after the prime than before it. Not that the distribution of change in risk tolerance in the male-identity prime subsample displays a fat right tail. For the sake of robustness, In Table B.1 of Online Appendix B I show that the subjects in this tails do not drive the results.

In Panel B of Table 3, I add the subsample of women to the analysis, and hence I use a triple-differences estimator. I estimate the following:

$$\begin{aligned}\Delta Risk\ Tolerance_{iae} = \alpha + \beta Man_{iae} + \gamma_1 \times Male\ Prime_{iae} + \gamma_2 \times Male\ Threat_{iae} \\ \gamma_3 \times (Male\ Prime \times Man)_{iae} + \gamma_4 \times (Male\ Threat \times Man)_{iae} + \eta_a + \eta_e + \epsilon_{iae}, \quad (2)\end{aligned}$$

where $\Delta Risk\ Tolerance$, $Male\ Prime$, and $Male\ Threat$ are defined as in equation (1),

[23] A subject-level specification whose outcome variable is the change in the number of risky choices like equation (1) is of course equivalent to one in which the outcome variable is the number of risky choices and the running sample includes two observations per subject, one before and one after the experimental conditions, with interactions of a pre-post dummy on the right hand side (for instance, see D'Acunto et al. (2019d)).

and Man equals 1 if the subject is a man.

Note that in equation (2), β does *not* capture the difference in average risk tolerance between men and women before the experiment, but captures the difference in the *change* in risk tolerance due to the exposure to the control condition. Hence, $\hat{\beta}$ should not be different from zero economically or statistically, unless being exposed to the control condition affected men and women's risk tolerance differently. I cannot reject the null that $\hat{\beta}$ is zero at any meaningful level of significance.

The increase in risk-tolerant choices by men in the male-identity-prime condition is significantly more positive (0.84, s.e. 0.37) than that of women, and with respect to the control group. The result is robust to averaging out age- and education- group effects. Men also became more risk tolerant after the identity threat, for which I can reject the null of no effect at the 10% level of significance (0.59, s.e. 0.33).

3 Economic Channels: Subjective Beliefs and Non-gendered Norms

After having established that priming or threatening male identity increases men's willingness to take risk, I move on to test for the economic channels that might explain this effect. The primes may act through two channels: (i) Preferences: Priming identity could be a positive shock to men's risk tolerance. This shock would reduce the certainty equivalent men require to give up a chance to take part in a lottery. (ii) Subjective beliefs: Priming identity could be a positive shock to men's subjective probability of experiencing good outcomes if they take part in a lottery. They could believe they are more likely to obtain the good outcome conditional on participating than what the objective probabilities suggest.

Identity theory includes a direct role for the subjective-beliefs channel. Stets and Cast (2007) argue that "individuals manipulate beliefs as personal resources, including one's worth and efficacy. Positive and efficacious self-feelings help individuals be more persistent in the face of difficulties in verifying the self." Non-standard beliefs are motivated, because they have a self-serving scope—avoiding the incongruence of a lack of identity verification.

A Identity and Subjective Beliefs

Testing for an effect of identity manipulations on beliefs requires a procedure that does not simultaneously elicit risk attitudes and the subjective beliefs of experiencing good outcomes. I propose a novel procedures for this test in Experiment 2.

Experiment 2: Design and Procedure. I recruited 325 subjects on mTurk in May 2015 and invited them to work on a survey to earn \$.40, plus a bonus based on their answers. The bonus could range between \$0 and \$.55 based on subjects' answers and luck. I excluded two subjects (1% of sample) because they did not write meaningful words or sentences in the priming task, leading to a final sample of 323 subjects. After answering the same four background questions, subjects were exposed to the experimental conditions of Experiment 1 (control, male-identity prime, female-identity prime). I assigned subjects in equal proportions to the three experimental conditions (see Table 2).

I elicited subjective beliefs in two steps. First, subjects faced a lottery that won on average five out of 10 times. No additional details were disclosed, including the "winning" or "losing" payoffs. Subjects predicted the number of times they would win if they played the lottery 10 times. Asking subjects to think they were playing the lottery 10 times induced personal involvement, a direct action taken by subjects that motivates the arising of the illusion of controlling a random outcome (e.g., see Thompson (1999)). To make the elicitation incentive-compatible, subjects knew the lotteries would be played 10 times using a random-number generator after the experiment, and they would receive a bonus based on the difference between the predicted number of wins and the actual number of wins.[24] This design ensured that (i) subjects were told explicitly that the outcomes of the lotteries would be completely determined by chance, and (ii) risk tolerance had no role in the predictions, because subjects had no choice between the lottery and a certain outcome.

In a second step, subjects were told to think about one person at random they knew in their neighborhood, with no further detail. Subjects then predicted the number of times they thought this neighbor would win. This second part of the task is crucial to interpret subjects' beliefs in the first part. Under identity theory, subjects should overestimate the probability they would win in a pure game of chance due to motivated beliefs, but *not*

[24]The bonus was \$.10 if the subject predicted the actual number of wins, \$.05 cents if the subject was one time long or short the actual number of wins, \$.02 cents for two times long or short, and \$0.01 cent for three times long or short.

the probability anybody else would win, because there is no scope for motivated beliefs in others' actions. To the contrary, if subjects had overestimated both their winning probability as well as others' winning probability, either subjects had not understood the objective probability of winning or the prime induced a state of elation in subjects, which could have distorted their perception and beliefs (e.g., see Kuhnen and Chiao (2009)). The second part of the task is thus important to provide a falsification test for the emergence of motivated beliefs induced by identity norms.

Experiment 2: Results. Panel A of Figure 4 plots men's subjective beliefs about winning the lottery, across the control and primed conditions, where the primed condition pooled together identity salience and threat.[25] Primed men thought they would succeed more than five times on average, whereas control men did not. A double-sided t-test for whether the average predicted number of wins by primed men (5.25) equals five rejects the null at the 1% level of significance, whereas a similar test for whether the average number of wins by control men (4.96) equals five cannot reject the null at any plausible level. Testing if the averages across groups are equal rejects the null at the 10% level of significance.

Panel B of Figure 4 plots the average difference between men's predicted wins by themselves and by peers across experimental conditions. The difference (0.04) is statistically and economically indistinguishable from zero for controls. It is significantly positive for primed men (0.25). Manipulating identity therefore does increase the subjective probability of experiencing a winning lottery outcome for men, even in a pure game of chance. This result is consistent with overconfidence in the form of the illusion of controlling a game of chance.

In Figure B.4 of Online Appendix B, I show the manipulations had no effect on the beliefs of women. Both control and treated women thought on average they would succeed five times if they played the lottery 10 times, in line with the objective probabilities they were told. The different results on the effect of the manipulations on beliefs across men and women pair up with the asymmetric effect of identity manipulation on risk attitudes by gender in Experiment 1.

A potential alternative interpretation is that men know the "correct" answer but want to

[25] I merge the identity-salience and -threat groups to increase the power of the test, because the two primes have very similar effects on the willingness to take on risks in Experiment 1.

be ironic.[26] Ruling out this possibility with a formal test is hard, but three arguments reduce this concern. First, if men wanted to be ironic, it would be unclear why they reported beliefs that were systematically biased upwards, as opposed to reporting any number different from 50% above or below. Second, it would also be unclear why they would report distorted beliefs about their own probability of winning but objective beliefs about others' probability of winning. Third, the systematic heterogeneity of the baseline effect based on commitment to gender identity I discuss in Section 5 would imply that older men and men in the US South are more ironic than other men, for which I could not find any evidence in the psychology literature.

B Norms of Behavior vs. Gender: Priming Power Identity

A second corollary of identity theory is that the gender context in itself is not crucial for the effects of identity on men's choice under risk. In fact, *any* social identity to which men commit and that prescribes the same norms of behavior as male identity should have similar effects on their choice under risk. To assess this prediction, in Experiment 3, I follow the social psychology literature (e.g., see Anderson and Galinsky (2006); Fast et al. (2012)) and prime a powerful identity to subjects, by asking subjects to recall and describe in detail a situation in which they had power over one or more individuals. Power identity is a viable candidate, because Table 1 shows the stereotypes attached to powerful individuals are strikingly similar to those attached to men. Powerful identity thus should prime the same norms of behavior in men as male identity, even if it includes no reference to gender.[27]

Recalling a situation of power over others may also induce a state of elation in subjects, which would affect risk attitudes for reasons unrelated to identity norms (Kuhnen and Knutson (2011)). To assuage this concern, I prime the sense of success in subjects as a placebo manipulation. Recalling a situation in which subjects were successful and describing it in detail likely induces a state of elation in subjects. But whereas powerful individuals are described with norms of behavior similar to those of men, successful individuals are associated with a mix of male- and female-identity norms of behavior (see Table 1).

[26] I thank Kelly Shue for proposing this interpretation.

[27] Fast et al. (2012) interpret this test also as a way to induce the illusion of controlling random outcomes, which might thus also be interpreted as another test of the subjective-beliefs channel.

Experiment 3: Design and Procedure. I recruited 340 subjects on mTurk (mTurk) in May 2012 (first session) and October 2012 (second session), and invited them to work on a survey on creative writing to earn $.50, and a bonus by picking lotteries. I restricted the subject pool to mTurk Workers in the US and with more than 95% of lifetime tasks approved by Requesters. Table 2 describes the experimental design, which was a 2 (male, female) X 3 (control, power prime, success prime) factorial design. I excluded 17 subjects (5% of sample) who switched across the certainty equivalent and the lottery more than once, leading to a final sample of 323 subjects. I assigned subjects in equal proportions to the three experimental conditions.

The first and third stages of the experimental procedure (choosing lotteries) were exactly the same as in Experiment 1. In the second stage, subjects in the control condition recalled and described a situation in which they felt relaxed. Subjects in the overconfidence-prime condition recalled and described a situation in which they had power over an individual or individuals, whereas those in the success-prime condition a situation in which they felt successful.

Experiment 3: Results. I measure subjects' tolerance to risk before and after the treatment, as well as its change, as in Experiment 1. Panel A of Table 3 reports the estimated coefficients for the following OLS equation:

$$\Delta Risk\ Tolerance_{iae} = \alpha + \gamma_1 \times Power\ Prime_{iae} + \gamma_2 \times Success\ Prime_{iae} + \eta_a + \eta_e + \epsilon_{iae}. \quad (3)$$

Column (3) of Panel A of Table 3 shows that men primed with power over others chose lotteries over certainty equivalents 0.52 (s.e. 0.26) times more after the prime than before, relative to controls. The effect is robust to adding age- and education-group fixed effects.[28]

The size of the effect of the power prime on men's risk tolerance is similar to the size of the effect of the identity manipulations on risk tolerance in Experiment 1. This similarity is consistent with the similar characterizations of male and powerful individuals in Table 1. Once I add women in Panel B of Table 3, the effect of power on men is still positive and similar in estimated size (0.44), but I cannot reject the null that the coefficient is zero at plausible levels (s.e. 0.42). The estimated coefficient on the elation-state treatment, *Success Prime*, is always close to zero in magnitude and not statistically different from zero.

[28] One subject did not report his education level.

Contrary to Experiment 1, the effect of the control condition on the omitted categories is on average negative. A t-test for whether the mean is different from zero within the omitted categories cannot reject the null at any conventional level of significance. But the low power of a test with few observations might drive this non-result, which is why, in the next section that looks at risky decisions framed as investment opportunities, I asked *all* subjects to recall and describe a situation in which they felt relaxed, which primed the same state as the control manipulation in Experiment 1 and Experiment 3.[29]

4 Robustness and Applications

So far, I have tested for the causal effect of male identity norms on risk taking and the channels of transmission through stylized lottery-choice settings. In this section, I consider two applications to well-documented regularities in the literature, which suggests that men tend to take on more risk than women in the context of direct and delegated investment choices.[30] I design choices framed as investments and prime male and power identity in a between-subjects design. These tests can also be interpreted as a replication of the baseline results in Experiment 1 and an assessment of whether the framing of the tasks, setting, and context, rather than the priming of identity norms, drove the baseline results.

A Identity and Investment Choices

Why do agents invest in risky opportunities whose expected value is typically negative, such as individual stocks, even if they are risk averse? This behavior is hard to reconcile with expected utility theory. Identity theory can accommodate this fact: Men should make riskier choices on average irrespective of the opportunities' expected value, because their utility increases when they verify their identity. In Experiment 1.A, I test how identity affects the choices of men facing simple investment decisions with varying expected values.

Experiment 1.A: Design and Procedure. I recruited 240 subjects on mTurk in

[29]Experiment 3 also differs on the gender composition of the sample. About 60% of US mTurk Workers are women (Mason and Suri, 2012), but Experiment 3 has more men than women (see Table 2).

[30]For instance, see Eagly and Steffen (1984), Akerlof and Kranton (2000),Barber and Odean (2001), Niederle and Vesterlund (2007), Eckel and Grossman (2008), Croson and Gneezy (2009), Bertrand (2011), Dohmen et al. (2011), Mastripieri et al. (2009), Cesarini et al. (2010), Cronqvist and Siegel (2014), Cronqvist et al. (2014). See also Niederle on whether gender differences in risk attitudes are robust across elicitation methods and contexts.

September 2012 and invited them to work on a survey on creative writing for $.50, plus a bonus by performing simple investment decisions. I restricted the subject pool to mTurk workers in the US with more than 95% of lifetime tasks approved by Requesters. The experimental design was a 2 (male, female) X 3 (control, male-identity prime, power prime) factorial design. I excluded six subjects (3% of whole sample) who did not write meaningful words in the priming task, leading to a final sample of 234 subjects. I assigned subjects in equal proportions to the three experimental conditions.

The experimental procedure was as follows. In the first stage, subjects answered four background questions. Then *all* subjects recalled a situation in which they felt relaxed and wrote a short text describing this situation. This step differs from the procedure in Experiment 1, which had a within-subjects design. It is important, because in a between-subjects design I cannot test directly whether a reaction by control subjects drives the results. In the second stage, the control and male-prime conditions were the same as in Experiment 1. The power prime was the induction of a sense of power from Experiment 3.

The third stage consisted of the investment decisions. I gave subjects a virtual endowment of $100 at the beginning of each of three periods. Each period, subjects faced an opportunity and decided how much, if any, of their per-period endowment to invest (Gneezy and Potters (1997)). The (scaled) part of the endowment subjects did not invest, including potentially the whole endowment, was paid to them at the end of the experiment.

The three opportunities varied in terms of both expected value (EV) and riskiness. The first opportunity succeeded with probability 1/2 and paid off three times the invested amount in case of success (mid-risk, EV>0). The second and third opportunities both succeeded with probability 1/6. The second opportunity paid off 8 times the invested amount in case of success (high-risk, EV>0), whereas the third opportunity paid off only 5 times the invested amount in case of success with a probability of 1/6 (EV<0). The latter investment is a money-burning opportunity, because in expectations subjects would be better off by keeping their full endowment and investing nothing.

I presented the three opportunities to subjects in random order. Subjects could not invest less than zero, that is, pay to avoid a choice. A random-number generator calibrated to the objective probabilities of each opportunity determined good or bad outcomes

and subjects knew about this random-number-generator procedure. Subjects received feedback about the outcomes directly after their choices to test if the effects of identity and overconfidence were not confined to experiences of good outcomes but survived even after experiencing negative outcomes, which is what I find. I paid subjects a scaled amount of what they were left with at the end of the experiment in the form of a bonus payment on top of the \$0.5 show-up fee.

Experiment 1.A: Results. I estimate the following probit model to assess the effects of identity on the extensive margin of investment (whether to invest or not):

$$\Pr(Invest = 1)_{iae} = \Phi(\alpha + \gamma_1 \times Male\ Prime_{iae} + \gamma_2 \times Power\ Prime_{iae} + \eta_a + \eta_e), \quad (4)$$

where $Invest_{iae}$ equals 1 if subject i in age group a and education-level group e agreed to invest any positive amount, and zero otherwise. $\Phi(.)$ is the normal cdf, and covariates are defined as in equation (1). I cluster standard errors at the subject level, because subjects made three investment choices in this experiment and each subject-choice pair is a separate observation.

Panel A of Table 4 reports the average marginal effects computed at the mean of variables. Column (1) shows men primed with male identity were 12 percentage points (s.e. 6.5 p.p.) more likely to invest than controls, who invested on average 65% of the time. The effect is robust to averaging out fixed age and education effects. Men primed with power identity invested 17 percentage points (s.e. 6.8 p.p.) more often than controls.

Panel B of Table 4 refers to the intensive margin of investment (invested amount conditional on investing). The invested amount is censored to the right (\$100) and to the left (\$0); hence, I estimate a tobit specification. Covariates are as in equation (1). Because estimated coefficients are within the censoring interval, I interpret them similarly to OLS coefficients. Column (1) shows subjects primed with male identity invested on average \$16 dollars more in each opportunity than controls, but the effect is not statistically significant. Averaging out age and education effects (column (2)) increases the size of the estimated coefficient to about \$21 (s.e. \$10), which is an economically and statistically significant effect. Subjects primed with power invested about \$17 (s.e. \$9) more than control subjects in the baseline specification of column (1). The size of the effect increases to about \$23 (s.e. \$10) once age and education fixed effects are averaged out.

I then use Experiment 1.A to assess whether, consistent with identity theory, identity salience has similar effects on risk-taking irrespective of the quality of the investment opportunities men face. Panel A and Panel B of Figure 5 plot the results across opportunities. Qualitatively, the effect of identity is larger for riskier opportunities and twice as large for the risky negative-EV opportunity relative to the less risky profitable opportunity. The economic significance of the differences is not negligible, although they are not always statistically significant.

B Identity and Delegated Investment

As a second application, I consider delegated investment. Excessive risk-taking in this context is often explained through agency theory: Agents take on excessive risk when their incentives and the principal's are not aligned. Beliefs-based interpretation have also been proposed, such as the overconfidence hypothesis,[31] which suggests agents might take on excessive risk because they believe their excessive-risk choices are optimal for their principals. Identity theory provides an additional explanation: Agents might take on excessive delegated risk to confirm their own identity. Identity should thus affect men's choice similarly irrespective of whether the choice is direct or delegated, which is not true under an agency interpretation.

To test for the effects of identity on delegated choice under risk, I design a setting in which the interests of the principal and the agent are aligned perfectly, so that agency-based interpretations of any effects are muted. Specifically, the compensation of the agent is a fixed proportion of the principal's revenues. Incentives in the delegated choice and direct choice are the same and the two choices differ only in relation to their framing. Moreover, the setting is designed so that managers know the objective probabilities of good states if they choose a risky bet and there is no role for agent's ability in shaping the size of the payoffs across future states of the world. This feature of the design ensures that overconfidence-based interpretations cannot explain any different reaction by treated and control men.

The drawback of aligning principal and agent's interests perfectly is that subjects face choices that are framed as delegated, but in which the principal has not active role.

[31] See Malmendier and Tate (2005), Malmendier and Tate (2008), Hirshleifer, Low, and Teoh (2012), and Ben David, Graham, and Harvey (2013).

Unfortunately, the experimental setting does not allow me to create actual delegated choices, but a substantial body of research in behavioral economics and social psychology, starting with Tversky and Kahneman (1981), shows that framing has a substantial impact on choices in otherwise similar settings.

Experiment 1.B: Design and Procedure. I recruited 220 subjects on mTurk in October 2012, and invited them to work on a survey on creative writing for $.50, plus a chance to earn a bonus by performing simple investment decisions. I restricted the subject pool to Workers in the US who had more than 95% of tasks approved by Requesters. The experimental design was a 2 (male, female) X 2 (control, male identity prime) factorial design.[32] I excluded seven subjects (3% of whole sample) who did not write meaningful words in the priming task, leading to a final sample of 213 subjects. I assigned subjects in equal proportions to the experimental conditions.

The procedure was exactly the same as in Experiment 1.A in the first and second stage. That is, first subjects answered four background questions and then all subjects recalled a situation in which they felt relaxed and wrote a short text describing this situation. Then, in the second stage, the same control and male-prime conditions as in Experiment 1 and Experiment 3 were proposed randomly to subjects.

Experiment 1.B differed from Experiment 1.A in the third stage, in which investments were framed as delegated decisions. In each of two periods, subjects had to imagine they were given money by Sally, a lady in their neighborhood who wanted to invest her money instead of keeping it in her checking account. She had no clue of how to invest and was asking for subjects' help in managing her endowment. Each period, subjects faced an investment opportunity. They decided if they wanted to invest any of Sally's endowment, and if yes, how much. Their compensation would be a performance-based fee, calculated as a fraction of the amount Sally ended up with after all the investment decisions. Sally was the fictitious principal in this case and subjects acted as her agent.

The two opportunities had the same probability of succeeding but different EVs. One opportunity succeeded with probability 1/2 and paid off 2.2 times the invested amount (EV>0). The other opportunity succeeded with probability 1/2 and paid off 1.8 times the invested amount (EV<0). A risk-neutral agent would have invested her whole

[32] I have only considered the male identity prime to increase the power of the test, given the strikingly similar effects of the identity and power primes in Experiment 1.A.

endowment in the first case, and nothing in the second case. The opportunities were presented in random order. Each period, subjects could not invest more than their endowment or less than zero.

Experiment 1.B: Results. Columns (3) and (4) of Panel A of Table 4 report the marginal effects implied by the estimated coefficients of equation (4). Men whose identity was primed were 9.7 percentage points more likely than control men to invest, but I cannot reject the null that this coefficient equals zero statistically (s.e. 6.9 p.p.). The effect is similarly positive but statistically insignificant when controlling for age and education effects, as well as when computing the marginal effects with a linear probability model (untabulated).

Panel B of Table 4 shows results for estimating a tobit specification whose dependent variable is the invested amount in the two opportunities, which is censored at \$0 and \$100. Men whose identity was primed invested on average \$22.9 more (s.e. \$10.38) than controls. The effect is robust to controlling for age and education effects, and is not dissimilar in size from the effects on direct investment opportunities (columns (1) and (2)).

Overall, men whose identity is salient do not seem to behave differently when facing direct or delegated investment opportunities, which is consistent with identity theory. In Table B.2 of Online Appendix B, I report the results of Experiment 1.A and 1.B for the female subsample, which, consistent with the lack of effects on risk tolerance in Experiment 1, does not appear to react to either female-identity prime or threat

5 Changing Identity Norms: Cohorts and Cultural Regions

Armed with the results so far, I can finally tackle the question that opened the paper: Could blurring gender identity norms have implications for men's economic decision-making under risk? According to Burke (1980b), the effects of identity on risk attitudes should be higher for men who commit more to male identity, because more committed men derive more utility when confirming their identity.

The main challenge in designing such test is that I cannot elicit directly within the experiment the extent to which subjects commit to male-identity norms. If I elicited commitment before the experimental manipulations, I would prime all subjects with identity norms. If I elicited commitment after the experiment, such as in a debriefing

session, I would not know whether the reported commitment was conditional or unconditional to the treatment effect of identity salience. Indirect methods to elicit subjects' commitment to social norms, such as the Implicit Association Test (see Carlana (2019) for a recent combination with field data), might still potentially prime all subjects with gender identity norms, although more subtly than a direct question.

Because of this challenge, I use two non-experimental sources of variation based on the varied locations and demographics of my subject pool. Using non-experimental variation also helps in terms of reducing concerns about demand effects, because it would be hard for subjects to fine tune their choices based on their (unknown) position in the distribution by commitment to identity (see also Section 6).

In the first test, I compare the reactions of men who belong to different cohorts. The rationale is that commitment to gender-identity norms has decreased substantially over the last few decades.[33] Moreover, the use of masculine individuating cues increases with age, and the dichotomy between masculine and feminine stereotypes also becomes starker with age (Biernat (1991)). Men in older cohorts should thus be more committed to male identity relative to men in younger cohorts, because of both a cohort and an age effect.

In the second test, I exploit variation in identity commitment across space by comparing the reactions of men who live in the Southern US to the reactions of other men. The rationale for this test is that men in the US South tend to have more conservative views about gender roles than men in the rest of the US, based on survey evidence such as the answers to questions in the General Social Survey (e.g., Rice and Coates (1995); Powers et al. (2003); Davis and Greenstein (2009)).

Although using non-experimental variation has the advantages I describe above, it also includes a set of caveats. First and foremost, I could not stratify the experimental samples based on these two sources of variation—cohorts and state-level geography—and hence I could not design the experiments to maximize statistical power and representativeness across cells. To reduce this concern, I merge the treated experimental cells in experiments with two treatment conditions for both of which identity theory has the same predictions, such as the male identity and power identity priming conditions in Experiment 3.

[33]For example, see Ciabattari (2001), Fernandez, Fogli, and Olivetti (2004), Goldin et al. (2006), Doepke and Tertilt (2009), Alesina, Giuliano, and Nunn (2013), Glaeser and Ma (2013), Breda and Ly (2014), and Olivetti, Patacchini, and Zenou (2015).

A second caveat is that the relative sizes of split samples are not always consistent across experiments. This issue is especially evident for the sample split across space. For instance, in Experiment 1, Experiment 2, and Experiment 3, the number of subjects from Southern states is about one half the number of subjects from other states, whereas in Experiment 1.A and Experiment 1.B, this proportion reverses. This discrepancy might relate to the timing of the launch of the experiments. The two last experiments started to recruit subjects around 5PM EST, whereas the other experiments started around 10AM EST. Accesses on mTurk might be increasing after close of business when mTurk Workers leave their regular jobs, and hence the last two experiments might oversample subjects leaving in the Eastern US compared to the Western US, which would increase the proportion of Southern subjects. Ultimately, though, I cannot explain the patterns of subject recruiting across space, because I could not stratify based on space.

I report the results for the effects across cohorts in Panel A of Table 5. "Young" includes subjects who were between 18 and 28 years old at the time of the experiment, and "Not Young" all the other subjects. Treated men in older cohorts increased their risk attitudes (column (1) and column (3)), displayed motivated beliefs in a pure game of chance (column (2)), and were more likely to invest and to invest larger amounts (columns (4)-(8)) relative to control men in their same age group, both economically and statistically. Instead, the reactions of treated men in the young cohort relative to the control men in their same cohort were substantially smaller in terms of economic magnitudes, and I cannot reject the null hypothesis of no difference. Statistical significance is sparser if I estimate (untabulated) triple interactions in specifications similar to the ones in the previous sections.

Overall, I interpret the results in Panel A of Table 5 as broadly consistent with the notion that identity has stronger effects on older cohorts of men.

In Figure 6, I further assess the shape of the reaction across different cohorts in graphical form. I report the analyses of Experiment 1, Experiment 1.A, and Experiment 1.B separately for four cohorts of men—subjects born before 1975, between 1976 and 1983, between 1984 and 1989, and between 1990 and 1994.[34] Panel A of Figure 6 shows the effect of identity salience on the change in risk tolerance is highest for subjects born before 1975, and it declines monotonically to become insignificantly negative for men

[34]These cohort boundaries were effectively set exogenously in the initial demographic questionnaire filled by all subjects.

born after 1990. A similar pattern exists for the effect of male identity on the probability of investing (Panel B of Figure 6) and for the amount invested.[35] Although statistical significance is sparse if I split the male sample into four groups, I again interpret these results as broadly consistent with the notion that identity affect the choice under risk of older men more than that of younger men.

The results across cohorts also speak to the external validity of the effects in the paper. Older individuals are on average wealthier and more likely to invest for retirement and in risky opportunities than the younger cohorts. The effects of identity salience on choice under risk are thus higher for the demographic group that is most relevant for such choices in the field.

I do not aim to disentangle the age effect from the cohort effect, because identity theory has the same predictions for both dimensions. The two mechanisms though have vastly different implications for the evolution of identity effects over time. If cohort effects drove the results, the aggregate effect of identity on risk attitudes and choice should fade over time as more individuals with less stark views on gender roles replace the older cohorts (e.g., see Guiso et al. (2008) for a cross-country argument). If aging partly drove the results, instead, the aggregate effects would decline more slowly over time, because men would become more reactive to identity as they aged.

One might be concerned the results capture some form of regression to the mean. Young men are more likely to engage in risky behaviors than older men. Young people may be so risk tolerant that the manipulations have no effect on them. But the average risk tolerance of men before the manipulations, the average likelihood of investing, and the average amounts invested by controls do not decrease with age, which is what the regression-to-the-mean explanation requires.

Moving on to the effects across space, I define subjects who connected to mTurk from IP addresses located in a Southern US state as coming from the "US South" and all other subjects as coming from the "Non US South." The US South includes the following US

[35] In this case, the pattern is decreasing but not monotonic

Census Bureau divisions: South Atlantic, East South Central, and West South Central.[36]

I report the results in Panel B of Table 5. Although I cannot always reject the null that the differences in choice and beliefs between treated and control men in the South are statistically different from zero, the sizes of these difference are larger than the sizes of the corresponding differences for men in non-Southern states. For non-Southern-state men, I always fail to reject the null that treated subjects make riskier choices relative to control subjects even when the sample size is substantially larger than the sample size of southern men.

6 Are Demand Effects a Relevant Concern?

Within-subjects designs combined with priming techniques contribute to raise concerns about demand effects. Subjects may be willing to conform their behavior to what they believe is the aim of the experimenter (see e.g., de Quidt, Vesterlund, and Wilson (2018b) and de Quidt, Haushofer, and Roth (2018a)).

As discussed above, though, identity theory predicts unambiguously that men should take on more risk after both priming and threatening their identity. Instead, under demand effects, men should take more risk after being primed with male identity, but should take less risk after being primed with female identity, because female identity prescribes risk-averse norms of behavior: Male subjects might think that after being primed with female identity they should behave more "feminine" instead of the opposite.

Experimenter demand effects also seem unlikely because men but not women react to the identity manipulations, and it is unclear why men but not women would be prone to demand effects. The cross-sectional results based on subjects' cohorts and geographic location are also barely consistent with demand effects. Demand effects could only drive these results if subjects were able to fine tune their choices to the unknown predictions about cohorts—without knowing the age distribution of the overall subject pool—as well as to the predictions about US states. That is, men would need to guess what is the age and geographic distribution of the rest of the subjects pool, guess the extent of

[36]These divisions include the following states: Alabama, Arkansas, Delaware, Florida, Georgia, Kentucky, Louisiana, Maryland, Mississippi, Oklahoma, North Carolina, South Carolina, Tennessee, Texas, Virginia, and West Virginia. Non-state territories include the District of Columbia. If I consider an alternative classification in which Delaware, Maryland, and the District of Columbia do not belong to the South, the results are virtually unchanged, because I have few subjects from Delaware or the District of Columbia across the five experiments.

risk taking of others, and place themselves in between according to their own age and location. And, all this should happen within the 15-minute time frame subjects take on average to complete the experiments, which seems rather implausible.

To assess the scope of demand effects more directly, I also propose a follow-up test.[37] In June and July 2014, I recruited 100 subjects not involved in any other experiments to perform the same tasks as in Experiment 1, but also to guess the motivations of the tasks at the end of the experiment. Guesses were incentivized, because subjects were paid a bonus if they guessed the correct motivation. Subjects answered two questions: "Why do you think we asked you to read a text and produce a short essay?" and "Why do you think we had you choose lotteries before and after the writing task?" The second question aimed to direct subjects explicitly to think about the connections between the priming procedure and the experimental tasks.

Figure 7 plots the reported motivations for the lottery tasks and the writing task across six categories. Only nine subjects in the "Change in Risk" category in Panel B of Figure 7 gave answers compatible with the true scope of the study, whereas, as I discuss in detail below, the extent of risky choices of virtually all men exposed to the treatment conditions is higher after the manipulation relative to controls. Of the nine subjects in the "Change in Risk" group, five were women and four were men. The men were all in the male-identity-prime condition; three of them were more risk tolerant after the manipulations, and one was not. Even if the subjects are not those in Experiment 1, the results suggest demand effects can hardly drive the results.

The generic categories "Relate to Lotteries" and "Relate to Text" include subjects who do not mention any scope consistent with the aims of the experiment. For instance, the answers of a man in the identity-threat condition are "Perhaps to get into a focused mindset" and "Maybe it was a part of a research and somehow connected." These categories also include two subjects who guessed the *opposite* direction of the effects predicted by identity theory: a man in the control condition ("I think you wanted to see if the regimented activity made people take less risk" and a man in the identity-threat condition ("It was a break from the lottery exercise, to re-orient us to feminine and less risky decision making"). Interesting examples are for men who thought the recall task was a distraction, and the aim was to test for the consistency of choices, but instead

[37] I thank Dan Benjamin for suggesting this test.

reacted to the manipulations. For instance, a man in the identity-salience condition replied, "This was to see how cautious I was in the bets despite your distraction." But then he made more risk-tolerant choices after the manipulation than before.

A Evidence against Fabricated Data

Finally, I run the test proposed by Simonsohn (2013) against fabricated data in experimental research. This test is important in light of the recent research-misconduct scandals commented on by Kahneman (2012). In Table B.3 of Online Appendix B, I verify for all the experiments in the paper that the standard deviations of the sample averages across experimental conditions do not differ from those one would obtain in random samples across conditions. In Experiment 1, the null that the observed distributions come from independently drawn random samples can only be rejected above the 99.99% significance level (99.98% in the male subsamples).

7 Discussion and Conclusions

A defining feature of identity theory rests on its welfare implications and hence on potential policy prescriptions. Typically, deviations from expected utility theory imply a drop in agents' utility, which justifies policy intervention to "rectify" the mistaken choice. The fact that identity verification increases utility, instead, suggests agents might gain from departing from expected utility theory even if they lose in terms of consumption utility. If the gains from identity verification overcame the consumption losses, the identity-based view would suggest that nudging agents to align their choices with the predictions of expected utility theory is ultimately detrimental to them. A major challenge for future research is providing reliable ways to compare consumption utility and the psychological utility deriving from identity verification so as to allow welfare statements and derive policy implications (e.g., see Guthrie and Sokolowsky (2013)).

A field application of the effects of gender identity norms on risk taking is the possibility for investment banks, financial advisers, and trading platforms to foster risk-taking in men and hence increase their revenues in terms of trading fees by making identity salient to clients. Anecdotal evidence consistent with this possibility is in Figure B.5 of Online Appendix B, which documents the use of explicit male-identity cues such as power, confidence, aggressiveness, and masculinity in the marketing of financial products. The

fact that cues often unrelated to the actual financial decision at stake do affect financial decision-making has been established in the literature (e.g., see Ravina (2008)). For the case of identity, investigating how the supply-side of financial services uses such cues to maximize the surplus they can extract from investors is an interesting direction for future research.

Another broad open area is the potential interaction between identity, biology, genetics, and neurology in determining the risk attitudes and choices of men and women. Despite the advantages of online platforms, such platforms do not allow reliable collection of information on the biological and neurological characteristics of subjects. Field settings the experimenter can physically reach may allow progress in this direction.

References

G. Akerlof and R. Kranton. Economics and identity. *Quarterly Journal of Economics*, 3 (115):715–753, 2000.

A. Alesina, P. Giuliano, and N. Nunn. On the origins of gender roles: Women and the plough. *Quarterly Journal of Economics*, 2(128):469–530, 2013.

C. Anderson and A. Galinsky. Power, optimism, and the proclivity for risk. *European Journal of Social Psychology*, 36:511–536, 2006.

J. Andreoni and C. Sprenger. Estimating time preferences from convex budgets. *American Economic Review*, 102(7):3333–3356, 2012.

B. Barber and T. Odean. Gender, overconfidence, and common stock trading. *Quarterly Journal of Economics*, 1(116):261–292, 2001.

G. Barone, F. D'Acunto, and G. Narciso. Telecracy: Testing for channels of persuasion. *American Economic Journal: Economic Policy*, 7(2):30–60, 2015.

I Ben David, J. Graham, and C. Harvey. Managerial miscalibration. *Quarterly Journal of Economics*, 4(128):1547–1584, 2013.

R. Benabou. The economics of motivated beliefs. *Revue d'économie poitique*, 125(5): 665–685, 2015.

D. Benjamin, J. Choi, and J. Strickland. Social identity and preferences. *American Economic Review*, 3(100):1913–1928, 2010.

D. Benjamin, J. Choi, and G. Fisher. Religious identity and economic behavior. *Working Paper*, 2013.

A. Berinsky, G. Huber, and G. Lenz. Evaluating online labor markets for experimental research: Amazon.com's mechanical turk. *Political Analysis*, 3(20):351–368, 2012.

M. Bertrand. New perspectives on gender. *Handbook of Labor Economics*, 4(B):1543–1590, 2011.

M. Biernat. Gender stereotypes and the relationship between masculinity and femininity: A developmental analysis. *Journal of Personality and Social Psychology*, 61(3):351–365, 1991.

P. Bordalo, N. Gennaioli, and A. Shleifer. Salience theory of choice under risk. *Quarterly Journal of Economics*, (127(3)):1243–1285, 2012.

P. Bordalo, K. Coffman, N. Gennaioli, and A. Shleifer. Beliefs about gender. *Working Paper*, 2018.

T. Breda and S. Ly. Stereotypes, discrimination, and the gender gap in science. *Working Paper*, 2014.

P.J. Burke. The self: Measurement implications from a symbolic interactionist perspective. *Social Psychology Quarterly*, 43:18–29, 1980a.

P.J. Burke. Identity processes and social stress. *American Sociological Review*, 56(6): 836–849, 1980b.
M. Carlana. Implicit stereotypes: Evidence from teachers' gender bias. *Quarterly Journal of Economics*, 3(134):1163–1224, 2019.
L. Carli. Gender differences in interaction style and influence. *Journal of Personality and Social Psychology*, 56:565–576, 1989.
K. Casler, L. Bickel, and E. Hackett. Separate but equal? a comparison of participants and data gathered via amazon's mturk, social media, and face-to-face behavioral testing. *Computer in Human Behavior*, pages 2156–2160, 2013.
D. Cesarini, M. Johannesson, P. Lichtenstein, O. Sandewall, and B. Wallace. Genetic variation in financial decision making. *Journal of Finance*, (65):1725–1754, 2010.
T. Ciabattari. Changes in men's conservative gender ideologies. *Gender & Society*, 15(4): 574–591, 2001.
K. Coffman. Evidence on self-stereotyping and the contribution of ideas. *Quarterly Journal of Economics*, (129):1625–1660, 2014.
A. Cohn, J. Engelmann, E. Fehr, and M.A. Marechal. Evidence for countercyclical risk aversion: An experiment with financial professionals. *American Economic Review*, 2 (105):860–885, 2015.
H. Cronqvist and S. Siegel. The genetics of investment biases. *Journal of Financial Economics*, 2014.
H. Cronqvist, A. Previtero, S. Siegel, and R. White. The fetal origin hypothesis in finance: Prenatal environment and financial risk taking. *Working paper*, 2014.
R. Croson and U. Gneezy. Gender differences in preferences. *Journal of Economic Literature*, 2(47):1–17, 2009.
F. D'Acunto. Tear down this wall street: Anti-finance rhetoric, motivated beliefs, and investment. *Working Paper*, 2018.
F. D'Acunto, M. Prokopczuk, and M. Weber. Historical antisemitism, ethnic specialization, and financial development. *Review of Economic Studies, forthcoming*, 2018.
F. D'Acunto, D. Hoang, M. Paloviita, and M. Weber. Cognitive abilities and inflation expectations. In *AEA Papers and Proceedings*, volume 109, pages 562–66, 2019a.
F. D'Acunto, D. Hoang, M. Paloviita, and M. Weber. Iq, expectations, and choice. Technical report, National Bureau of Economic Research, 2019b.
F. D'Acunto, U. Malmendier, and M. Weber. Gender roles distort women's economic outlook. *Working Paper*, 2019c.
F. D'Acunto, N. Prabhala, and A. Rossi. The promises and pitfalls of robo-advising. *The Review of Financial Studies*, 32(5):1983–2020, 2019d.
S. Davis and T. Greenstein. Gender ideology: Components, predictors, and consequences. *Annual Review of Sociology*, 35:87–105, 2009.
J. de Quidt, J. Haushofer, and C. Roth. Measuring and bounding experimenter demand. *American Economic Review*, 108(11):3266–3302, 2018a.
J. de Quidt, L. Vesterlund, and A. Wilson. Experimenter demand effects. *Handbook of Research Methods and Applications in Experimental Economics, forthcoming*, 2018b.
S. DellaVigna and D. Pope. Predicting experimental results: Who knows what? *Journal of Political Economy, forthcoming*, 2018.
M Doepke and M. Tertilt. Women's liberation: What's in it for men? *Quarterly Journal of Economics*, 4(124):1541–1591, 2009.
T. Dohmen, A. Falk, D. Huffman, U. Sunde, J. Schup, and G. Wagner. individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 3(9), 2011.
A. Eagly. Sex differences in social behavior: A social-role interpretation. *Hillsdale, NJ: Erlbaum*, 1987.
A. Eagly and V. Steffen. Gender stereotypes stem from the distribution of women and men into social roles. *Journal of Personality and Social Psychology*, 46:735–754, 1984.
C. Eckel and P. Grossman. Men, women and risk aversion: Experimental evidence. *Handbook of Experimental Economics Results*, Elsevier, 2008.
C. Exley and J. Kessler. Motivated errors. *Working Paper*, 2018.

N. Fast, N. Sivanathan, N. Mayer, and A. Galinsky. Power and overconfident decision making. *Organizational Behavior and Human Decision Processes*, 2(117):249–260, 2012.
B. Ferman, S. Fiorin, M. Kanz, and G. Rao. Status goods: Experimental evidence from platinum credit cards. *Quarterly Journal of Economics*, (133(3)):1561–1595, 2018.
R. Fernandez, A. Fogli, and C. Olivetti. Mothers and sons: Preference formation and female labor force dynamics. *Quarterly Journal of Economics*, (119(4)):1249–1299, 2004.
C. Frydman and M. Mormann. The role of salience and attention on choice under risk: An experimental investigation. *Working Paper*, 2018.
C. Frydman and B. Wang. The impact of salience on investor behavior: Evidence from a natural experiment. *Working Paper*, 2018.
A. Gargano and A. Rossi. Does it pay to pay attention? *Review of Financial Studies, forthcoming*, 2018.
E. Glaeser and Y. Ma. The supply of gender stereotypes and discriminatory beliefs. *NBER working paper no. 19109*, 2013.
U. Gneezy and J. Potters. An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics*, (112):631–645, 1997.
C. Goldin. A grand gender convergence: Its last chapter. *American Economic Review*, 4 (104):1091–1119, 2014.
C. Goldin, L. Katz, and I. Kuziemko. The homecoming of american college women: The reversal of the college gender gap. *Journal of Economic Perspectives*, 4(20):133–156, 2006.
L. Guiso, F. Monte, P. Sapienza, and L. Zingales. Culture, gender, and math. *Science*, 1 (320):1164–1165, 2008.
K. Guthrie and J. Sokolowsky. Characterizing the choice environment: Feeling constrained as an impediment to happiness. *International Journal of Happiness and Development*, 1(3):261–273, 2013.
D. Hirshleifer, A. Low, and S.H. Teoh. Are overconfident ceos better innovators? *Journal of Finance*, 4(67):1457–1498, 2012.
C. Holt and S. Laury. Risk aversion and incentive effects in lottery choices. *American Economic Review*, 4(92):1644–1655, 2002.
J Horton, D. Rand, and R. Zeckhauser. The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, (14):399–425, 2011.
C. Kuhnen and J. Chiao. Genetic determinants of financial risk-taking. *PlosONE*, 2(4), 2009.
C. Kuhnen and B. Knutson. The influence of affect on beliefs, preferences, and financial decisions. *Journal of Financial and Quantitative Analysis*, 3(46), 2011.
I. Kuziemko, M. Norton, E. Saez, and S. Stantcheva. How elastic are preferences for redistribution? evidence from randomized survey experiments. *American Economic Review*, 4(105):1478–1508, 2015.
A. Landier, Y Ma, and D. Thesmar. New experimental evidence on expectations formation. *Working Paper*, 2017.
C. Lian, Y Ma, and C. Wang. Low interest rates and risk taking: Evidence from individual investment decisions. *Review of Financial Studies, forthcoming*, 2018.
J. List. Why economists should conduct field experiments and 14 tips for pulling one off. *Journal of Economic Perspectives*, 3(25):3–16, 2011.
A. Maas, M. Cadinu, G. Guarnieri, and A. Grasselli. Sexual harassment under social identity threat: the computer harassment paradigm. *Journal of Personality and Social Psichology*, 5(85):853–870, 2003.
U. Malmendier and G. Tate. Ceo overconfidence and corporate investments. *Journal of Finance*, 60(6):2661–2700, 2005.
U. Malmendier and G. Tate. Who makes acquisitions? ceo overconfidence and the market's reaction. *Journal of Financial Economics*, 2008.
W. Mason and S. Suri. Conducting behavioral research on mechanical turk. *Behavioral Research Methods*, 44(1):1–23, 2012.
D. Mastripieri, P. Sapienza, and L. Zingales. Gender differences in financial risk aversion

and career choices are affected by testosterone. *Proceedings of the National Academy of Sciences*, pages 526–556, 2009.
G.H. Mead. Mind, self and society. *University of Chicago Press*, 1934.
K. Meier-Pesti and E. Penz. Masculinity and femininity as predictors of financial risk-taking: Evidence from a priming study on gender salience. *Journal of Economic Psychology*, 29:180–196, 2008.
P. Moore. The decline of the manly man. *YouGov Report*, 2016.
M. Niederle and L. Vesterlund. Do women shy away from competition? do men compete too much? *Quarterly Journal of Economics*, 122(3):1067–1101, 2007.
C. Olivetti, E. Patacchini, and Y. Zenou. Mothers, friends, and gender identity. *Working Paper*, 2015.
G. Paolacci, J. Chandler, and P. Ipeirotis. Running experiments on amazon mechanical turk. *Judgement and Decision Making*, 5(5):411–419, 2010.
R. Powers, Jill Suitor, and Susana et al. Guerra. Regional differences in gender. role attitudes: Variations by gender and race. *Gender Issues*, 2(21):40–54, 2003.
E. Ravina. Love & loans the effect of beauty and personal characteristics in credit markets. 2008.
E. Ravina. An assessment of the mating motive explanation of the beauty premium in market-based settings. *Behavioral and Brain Sciences*, 40, 2017.
T. Rice and D. Coates. Gender role attitudes in the southern united states. *Gender & Society*, 9(6):744–756, 1995.
C. Richards, W. Bouman, L. Seal, M. Barker, T. Nieder, and G. T'Sjoen. Non-binary or genderqueer genders. *International Review of Psychiatry*, (28(1)):95–102, 2016.
S. Schwartz and H. Tammy. Sex differences in value priorities: Cross-cultural and multimethod studies. *Journal of Personality and Social Psychology*, 89(6):1010–1028, 2005.
M. Shih, T. Pittinsky, and N. Ambady. Identity salience and shift in quantitative performance. *Psychological Science*, 10(1):80–83, 1999.
U. Simonsohn. Just post it: The lesson from two cases of fabricated data detected by statistics alone. *Psychological Science*, 24(10):1875–1888, 2013.
C. Steele. A threat in the air: how stereotypes shape intellectual identity and performance. *American Psychologist*, 52:613–629, 1997.
J. E. Stets and P.J. Burke. Identity theory and social identity theory. *Social Psychology Quarterly*, 63(3):224–237, 2000.
J. E. Stets and A.D. Cast. Resources and identity verification from an identity theory perspective. *Sociological Persepctives*, 50(3):517–543, 2007.
S. Stryker. Symbolic interactionism: A social structural version. *Blackburn Press*, 1980.
S. Thompson. Illusions of control: How we overestimate our personal influence. *Current Directions in Psychological Science*, 8(6):187–190, 1999.
A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981.
J. Vandello, J. Cohen, R. Burnaford, and J. Weaver. Precarious manhood. *Journal of Personality and Social Psychology*, 95(6):1325–1339, 2008.
R. Willer, C. Rogalin, B. Conlon, and M. Wojnowicz. Overdoing gender: a test of the masculine overcompensation hypothesis. *American Journal of Sociology*, 118:980–1022, 2013.

Figure 2: **Location of Subjects**

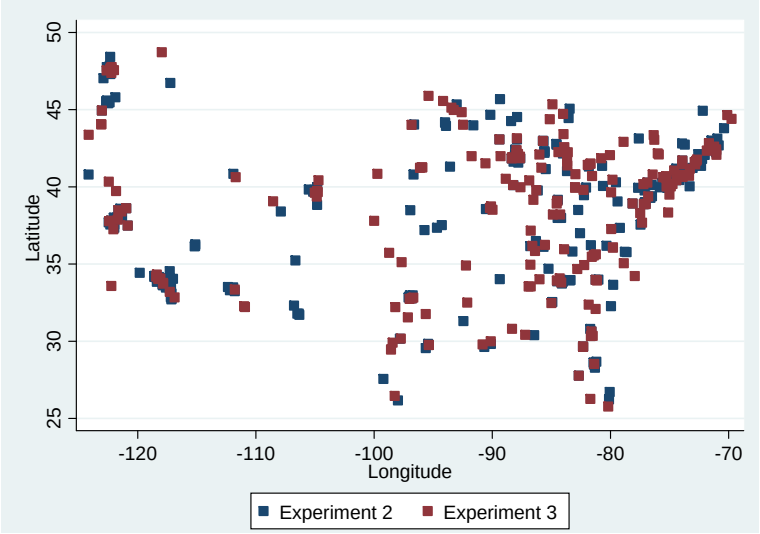
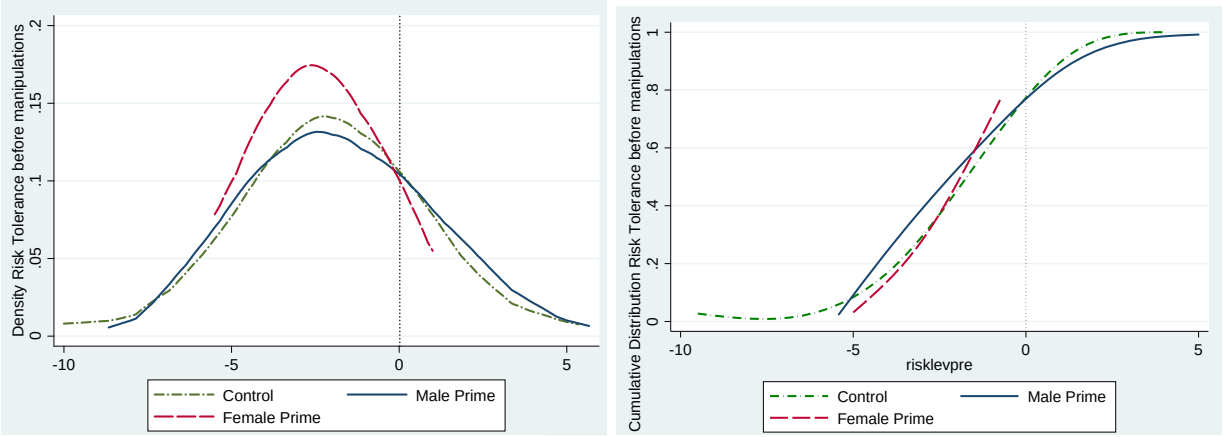


Figure 2 depicts the location of the subjects in Experiment 2 and Experiment 3 based on their IP address.

Figure 3: **Identity Salience and Changes in Risk Attitudes: Full Distribution**

(a) Distribution of Risk Tolerance before the manipulations (male subsample)



(b) Distribution of Δ Risk Tolerance induced by the manipulations (male subsample)

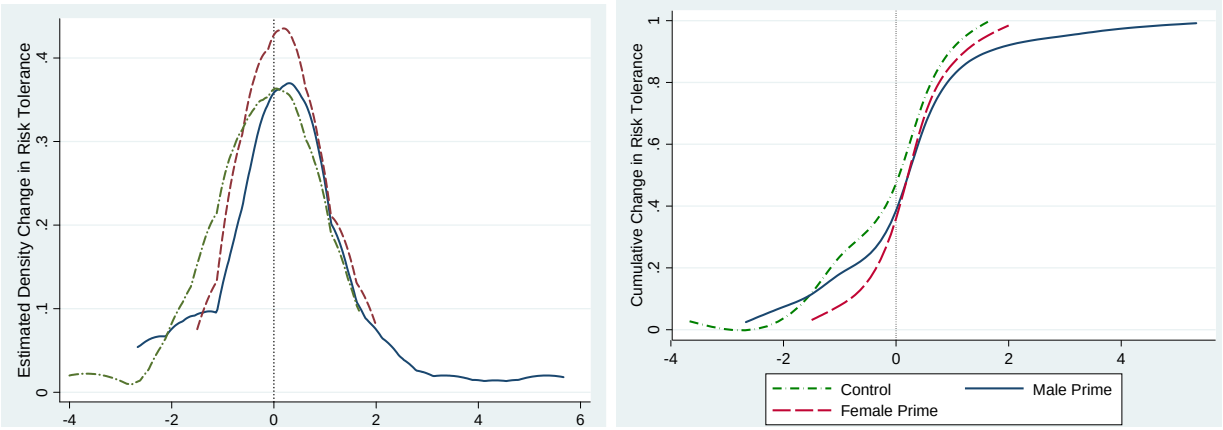
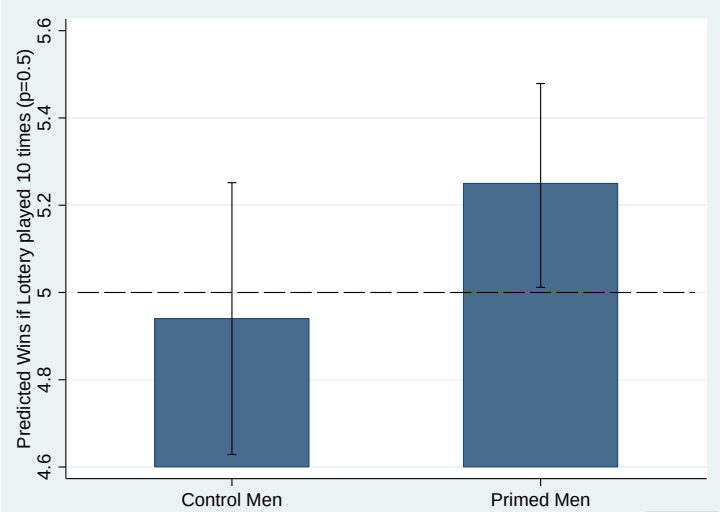


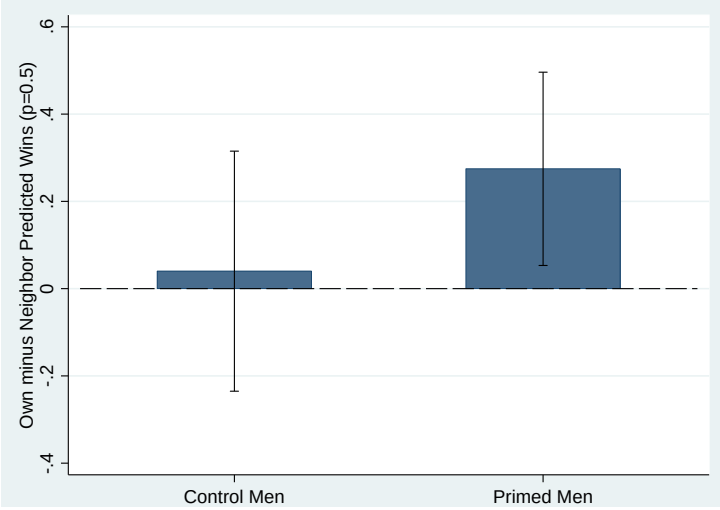
Figure 3 plots the estimated densities and cumulative distributions for the risk tolerance before the manipulations (panel A) and the subject-level change in risk tolerance before and after exposure to the experimental conditions (panel B) in Experiment 1 for male subjects. The subject-level change in risk tolerance is measured as follows: $\Delta Risk\ Tolerance_i = Risk\ Tolerance_{post,i} - Risk\ Tolerance_{pre,i}$, where $Risk\ Tolerance_{pre,i}$ and $Risk\ Tolerance_{post,i}$ are elicited via lottery choices a la Holt and Laury (2002). The smooth density plots use an Epanechnikov kernel with bandwidth of 0.7. The cumulative distribution plots interpolate the medians of 10 vertical bands with cubic splines. The plots of the distributions for the female subsample are in Figure B.4 of Online Appendix B.

Figure 4: **Identity Salience and Motivated Beliefs**

(a) Subjective beliefs of experiencing the good outcome in a lottery that succeeds on average 5 out of 10 times



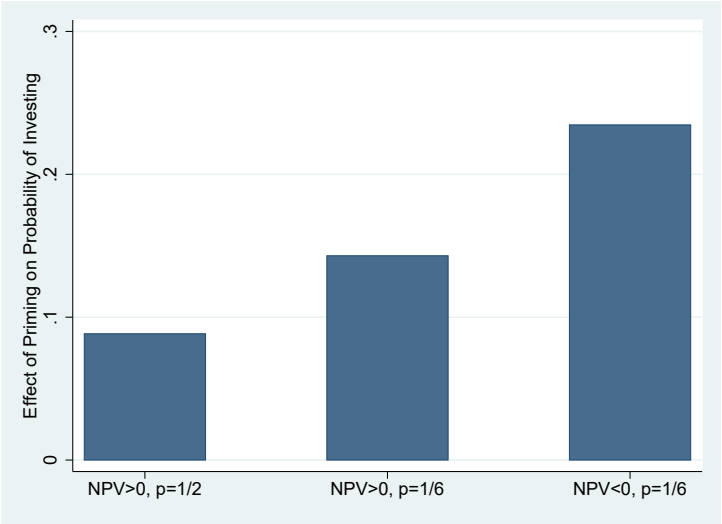
(b) Better-than-average beliefs across experimental conditions (men only)



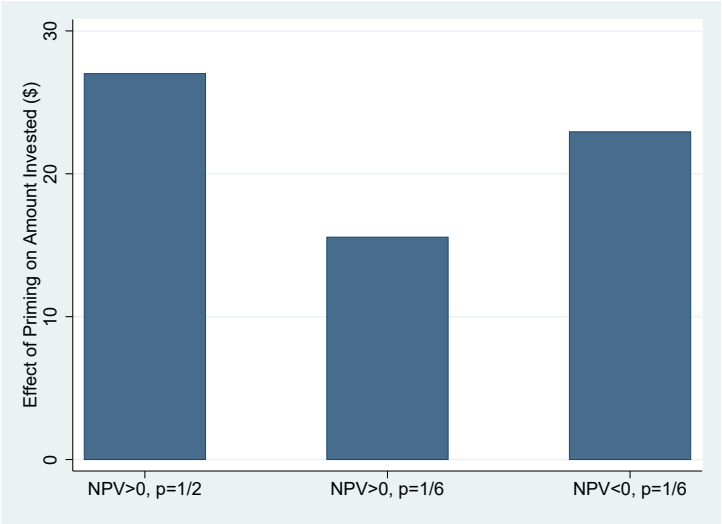
Panel (a) of Figure 4 plots the subjective beliefs of experiencing the good outcome by gender for subjects who imagine to play 10 times a lottery whose good outcome obtains on average 5 out of 10 times. Bandwidths are 95% confidence intervals for estimated means. Panel (b) of Figure 4 plots the average subjective better-than-average beliefs for men in the control group and in the primed group, which collects the male identity (salience) and female identity (threat) primes.

Figure 5: **Identity Salience and Direct Investment—Opportunities with Varying Expected Values (EV) and Risk Levels**

A. Effect of Identity Priming on the Probability of Investing by EV and Risk Levels



B. Effect of Identity Priming on the Amounts Invested by EV and Risk Levels



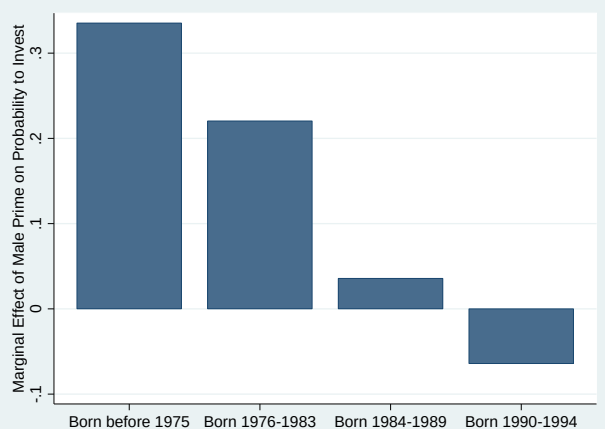
Panel A of Figure 5 reports the marginal effect attached to the treatment condition (male and power identity salience) when estimating the probit specification in Panel A of Table 4 in Experiment 1.A separately for each round of investment choice. Across the three rounds, subjects face a sound and moderate risk opportunity (EV>0, probability of winning=1/2), a sound and high risk opportunity (EV>0, probability of winning=1/6), and a money-burning and risky opportunity (EV<0, probability of winning=1/6). Panel B reports the marginal effect attached to the treatment condition (male and power identity salience) when estimating the tobit specification in Table 4 in Experiment 1.A separately for each round of the investment choices described above.

Figure 6: **Varying Commitment to Male Identity: Effects across Cohorts**

A. Effect of Male Identity prime on Risk Tolerance by male cohorts



B. Effect of Male Identity prime on Probability to Invest by male cohorts



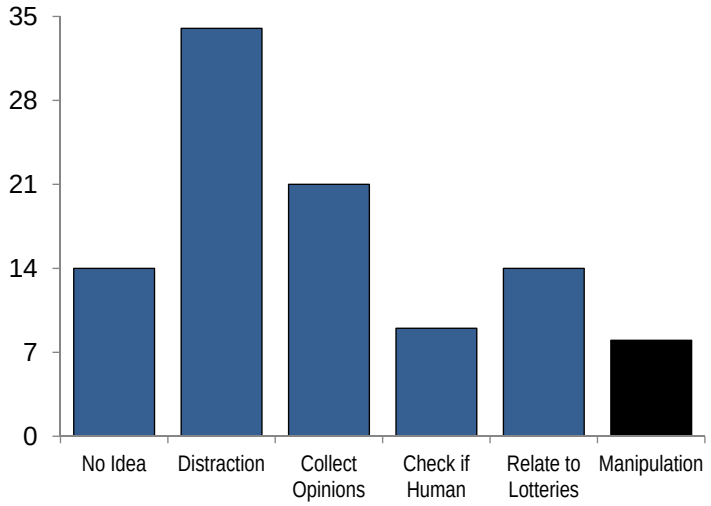
C. Effect of Male Identity prime on Amount Invested by male cohorts



Panels A-C of Figure 6 plot the effect of the male identity manipulations on the change in risk tolerance in Experiment 1, the likelihood that subjects invest in Experiment 1.A, and the average amount invested in Experiment 1.A, across four cohorts of subjects based on their year of birth.

Figure 7: **Testing for the Scope of Demand Effects**

A. Reported Motives for the Reading and Essay Tasks (N=100)



B. Reported Motives for the Lottery Tasks before and after Essay (N=100)

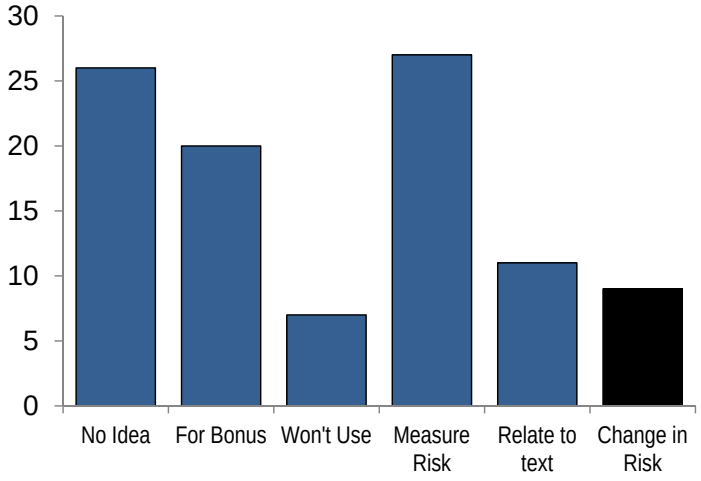


Figure 7 reports results for 100 subjects recruited in June and July 2014 to perform the tasks in Experiment 1 following the same recruitment procedures as described in the paper. These subjects were also asked to guess the motivations of the tasks at the end of the study. Panel A reports the distribution of answers for why subjects had to read a text and write a short essay, split into six broad categories. The dark histogram to the right (8 subjects) shows the answers compatible with the aim of the study. Panel B shows the distribution of answers for why subjects had to make lottery choices before and after the writing task. The emphasis on before and after aimed to instill the idea that the order of tasks was related to their motives. The only category compatible with the true motives includes 9 subjects (dark, right), 8 of which are those depicted in dark in Panel (a).

Table 1: **Normative Behaviors Attached to Identities by Online Users**

Male Identity		**Power Identity**		**Success Identity**		**Female Identity**	
Behavioral Norm	N	Behavioral Norm	N	Behavioral Norm	N	Behavioral Norm	N
Strong	37	Strong	50	Confident	46	Caring	40
Loud	27	Confident	45	Smart	36	Emotional	34
Aggressive	26	Arrogant	19	Happy	26	Flirty	16
Confident	18	Aggressive	18	Proud	19	Passive	15
Arrogant	15	Assertive	15	Determined	15	Smart	13
Dumb	12	Dominant	13	Focused	12	Loving	13
Stubborn	11	Greedy	9	Motivated	11	Happy	12
Assertive	11	Cocky	8	Assertive	10	Gossipy	12
Proud	10	Proud	8	Ambitious	10	Hepful	10
Dominant	10	Smart	7	Strong	10	Kind	10

Table 1 reports the words mentioned most frequently by 200 survey respondents recruited on mTurk – who did not act as subjects in any of the subsequent experiments – to describe how they think a typical male individual, powerful individual, successful individual, and female individual behave. Respondents wrote 3 to 5 words for each type of individual. Identities were presented in random order. Each normative behavior is followed by the number of times respondents reported that behavior or a synonim or verbal form.

Table 2: **Experimental Designs, Treatments, and Sample Sizes by Experimental Cells**

	Treatments	**Design**	**Number of Subjects**	**Main Results**
Experiment 1 Change Risk Attitudes	Male and Female Identity Prime	2 (male-M, female-F) X 3 (control-C, male identity-MI, female identity-FI)	MC:55, FC:74, MMI:61 FMI:64, MFI:32, FFI:34 Total: 320	Men whose identity is primed or threatened choose riskier lotteries after the primes. Women and control men do not react.
Experiment 2 Beliefs	Male and Female Identity Prime	2 (male-M, female-F) X 3 (control-C, male identity-MI, female identity-FI)	MC:50, FC:58, MMI:49 FMI:53, MFI:53, FFI:60 Total: 323	Men whose identity is primed or threatened have higher subjective beliefs of winning in a game of chance whose objective probabilities are known.
Experiment 3 Change Risk Attitudes	Power and Success Identity Prime	2 (male-M, female-F) X 3 (control-C, power identity-PI, success identity-SI	MC:55, FC:51, MPI:54 FPI:53, MSI:59, FSI:50 Total: 322	Men primed with power identity choose riskier lotteries after the prime. No effect of success primes (placebo) and no effects on women.
Experiment 1.A Negative EV Investment	Male and Power Identity Prime	2 (male-M, female-F) X 3 (control-C, male identity-MI, power identity-PI)	MC:35, FC:43, MMI:38 FMI:41, MPI:32, FPI:50 Total: 234	Men primed with male or power identity invest more, especially in negative expected value opportunities.
Experiment 1.B Delegated Investment	Male Identity Prime	2 (male-M, female-F) X 2 (control-C, male identity-MI)	MC:48, FC:60 MMI:49, FMI:55 Total: 212	Men primed with male identity invest more in delegated risky opportunities.

Table 3: **Identity and Risky Choices**

	(1)	(2)		(3)	(4)
	Experiment 1 Male and Female Identities			**Experiment 3** Power and Success Identities	
Outcome variable:	Δ *Risk Tolerance*			Δ *Risk Tolerance*	
	Panel A. Men Only				
Male Identity	0.507**	0.577**	Power Identity	0.504*	0.577**
	(0.246)	(0.264)		(0.261)	(0.273)
Female Identity	0.431**	0.564**	Success Identity	-0.002	-0.022
	(0.205)	(0.258)	(Placebo)	(0.262)	(0.263)
Mean Omitted	-0.212	-0.212	Mean Omitted	-0.473	-0.473
Median Omitted	0	0	Median Omitted	0	0
Session f.e.		X	Session f.e.		X
Age group f.e.		X	Age group f.e.		X
Education f.e.		X	Education f.e.		X
Observations	148	148	Observations	168	167
R-sq	0.03	0.09	R-sq	0.03	0.12
	Panel B. Men and Women				
Male Identity	-0.337	-0.286	Power Identity	0.077	0.124
	(0.280)	(0.268)		(0.331)	(0.348)
Man*	0.844**	0.836**	Man*	0.427	0.376
Male Identity	(0.373)	(0.375)	Power Identity	(0.422)	(0.446)
Female Identity	-0.163	-0.089	Success Identity	0.263	0.276
	(0.262)	(0.271)		(0.321)	(0.328)
Man*	0.594*	0.659*	Man*	-0.264	-0.310
Female Identity	(0.333)	(0.344)	Success Identity	(0.415)	(0.419)
Man	-0.0257	-0.329	Man	0.063	0.120
	(0.212)	(0.223)		(0.321)	(0.337)
Mean Omitted	0.045	0.045	Mean Omitted	-0.536	-0.536
Median Omitted	0	0	Median Omitted	-0.333	-0.333
Session f.e.		X	Session f.e.		X
Age group f.e.		X	Age group f.e.		X
Education f.e.		X	Education f.e.		X
Observations	320	320	Observations	322	321
R-sq	0.02	0.10	R-sq	0.02	0.08

Columns (1) and (2) of Table 3 report the estimated coefficients from:

$$\Delta Risk\ Tolerance_{iae} = \alpha + \gamma_1 \times Male\ Prime_{iae} + \gamma_2 \times Female\ Prime_{iae} + \eta_a + \eta_e + \epsilon_{iae}$$

Columns (3) and (4) report the estimated coefficient from:

$$\Delta Risk\ Tolerance_{iae} = \alpha + \gamma_1 \times Power\ Prime_{iae} + \gamma_2 \times Success\ Prime_{iae} + \eta_a + \eta_e + \epsilon_{iae},$$

where $\Delta Risk\ Tolerance_{iae}$ is the within-subject change in risk tolerance for subject i, in age group a and education group e, before and after being exposed to the treatments in Experiment 1 and Experiment 2. η_a and η_e are fixed effects for subjects' age and education-level groups. Hubert-White s.e. (equivalent to standard errors clustered at the subject level given the sample includes is one observation per subject) are reported in brackets. Statistical significance is reported as follows: ** 5%, * 10%.

Table 4: **Robustness & Applications: Direct and Delegated Investment**

	(1)	(2)		(3)	(4)
	Experiment 1.A Direct Investment			**Experiment 1.B** Delegated Investment	
Panel A. Probability of Investing (Probit Marginal Effects)					
Male Identity	0.123*	0.132**	Male Identity	0.097	0.071
	(0.065)	(0.063)		(0.069)	(0.066)
Power Identity	0.170**	0.179**			
	(0.068)	(0.070)			
Mean Omitted	0.65	0.65	Mean Omitted	0.84	0.84
Median Omitted	1	1	Median Omitted	1	1
Session f.e.		X	Session f.e.		X
Age group f.e.		X	Age group f.e.		X
Education f.e.		X	Education f.e.		X
Observations	315	315	Observations	194	194
Pseudo R-sq	0.03	0.09	Pseudo R-sq	0.01	0.06
Panel B. Amount Invested (Tobit Marginal Effects, limited between 0 and 100)					
Male Identity	15.66	21.09**	Male Identity	22.90**	20.93**
	(9.92)	(9.85)		(10.38)	(10.04)
Power Identity	16.97*	22.55**			
	(9.32)	(9.96)			
Mean Omitted	41.7	41.7	Mean Omitted	53.3	53.3
Median Omitted	29.0	29.0	Median Omitted	50.0	50.0
Session f.e.		X	Session f.e.		X
Age group f.e.		X	Age group f.e.		X
Education f.e.		X	Education f.e.		X
Observations	315	315	Observations	322	321
Pseudo R-sq	0.01	0.03	Pseudo R-sq	0.02	0.08

All the results in Table 4 refer to the subsample of male subjects. The corresponding (non-)results for the female subsample are reported in Table B.2 of Online Appendix B. Panel A of Table 4 reports the estimated coefficients from:

$$\Pr(Invest = 1)_{iae} = \alpha + \gamma_1 \times Male\ Prime_{iae} + \gamma_2 \times Power\ Prime_{iae} + \eta_a + \eta_e)$$

where η_a and η_e are fixed effects for subjects' age and education-level groups. Panel B of Table 4 reports the estimated coefficients from a tobit regression of the amount invested in each opportunity (censored at \$0 and \$100) on the Male identity prime and Power prime indicators from Experiment 1.A and Experiment 1.B. S.e. are clustered at the subject level. Statistical significance is marked as follows: ***1%, ** 5%, * 10%.

Table 5: **Changing Identity Norms: Effects by Cohorts and across Cultural Regions**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Experiment 1	**Experiment 2**	**Experiment 3**	**Experiment 1.A**		**Experiment 1.B**	
Outcome variable:	*Δ Risk Tolerance*	*Beliefs*	*Δ Risk Tolerance*	*Prob. Invest*	*Amount Invested*	*Prob. Invest*	*Amount Invested*
	Panel A. T-tests Treatment vs. Control Subjects, by Cohorts						
Young	0.438	0.167	0.128	-0.005	4.201	0.011	11.064
	(0.318)	(0.319)	(0.294)	(0.071)	(5.307)	(0.071)	(6.435)
Observations	67	63	92	166	166	126	126
Not Young	0.579**	0.460*	0.977***	0.325***	12.839**	0.232**	11.018
	(0.286)	(0.258)	(0.349)	(0.071)	(6.325)	(0.109)	(7.996)
Observations	81	87	72	147	147	68	68
	Panel B. T-tests Treatment vs. Control Subjects, by US Geographic Regions (non-missing location)						
U.S. South	1.363	0.665*	0.783*	0.152**	8.367	0.051	13.798**
	(0.840)	(0.337)	(0.419)	(0.063)	(5.170)	(0.078)	(6.513)
Observations	20	50	41	201	201	112	112
Not U.S. South	0.395	0.091	0.415	0.093	3.127	0.031	6.156
	(0.237)	(0.225)	(0.288)	(0.102)	(7.421)	(0.086)	(8.484)
Observations	52	99	115	87	87	62	62

Each cell of Table 5 reports the difference in the average value of the outcome variable reported on top of the table across men in the treatment and control group for each experiment across cohorts (Panel A) and across Southern and non-Southern United States (Panel B). Estimated Hubert-White standard errors reported in brackets below each coefficient refer to a two-sided t-test of the null hypothesis that the difference in the averages of the outcome variable across treated and control men in each cell equals zero.

Online Appendix

Identity and Choice Under Risk

Francesco D'Acunto, Boston College

Online Appendix A: Experimental Materials

This section includes the materials used in the experiments. Materials covering the same tasks were the same across experiments. Below, I list the materials by task and state in which experiments they were used:

- **Demographic Questions:** Experiments 1, 2, 3, 1.A, 1.B.
- **Experimental Manipulation Texts/Instructions:**
  - *Control*: Experiments 1, 2, 3, 1.A, 1.B;
  - *Male Identity*: Experiments 1, 2, 1.A, 1.B;
  - *Female Identity*: Experiments 1, 2;
  - *Power Identity*: Experiments 3, 1.A;
- **Lottery-Choice Task:** Experiments 1, 3;
- **Investment Tasks:**
  - *Direct Investment*: Experiment 1.A;
  - *Delegated Investment*: Experiments 1.B.

Demographic Questions: Experiments 1, 2, 3, 1.A, 1.B.

Country where you live

Age range

- 18-22
- 23-28
- 29-35
- 36-45
- 46-60
- 61+

Gender

- Male
- Female

Education

- No degree/Some High School (no degree)
- High School degree
- Some College (no degree)
- College degree
- Postgraduate studies

Experimental Manipulation Texts/Instructions

Control: Experiments 1, 2, 1.A, 1.B.

Dincharya [Daily Routine]

In Sanskrit, the word 'dincharya' means daily routine. According to Ayurveda, one should follow the dincharya in order to lead a healthy and disease-free life. Everyday, two cycles of change pass through the human body, each bringing a Vata, Pitta, or Kapha predominance. Based on the cycles of vata, pitta and kapha, our daily routine should be divided into morning, noon, evening/twilight, dinner and bedtime. In the Ayurvedic texts, it is written that a person should wake up two hours prior to the sunrise, if he/she is not suffering from any diseases such as fever or diarrhea. Very young, very old and sick people are some of the exceptions.

According to dincharya, the day should be kick-started by eliminating the colon and the bladder, followed by a through cleaning of the senses - ears, eyes, mouth etc. This should be followed by an oil self massage. Exercise in the morning, just after the massage, helps rejuvenate the body and soul. After bathing, one should head towards the dining table for breakfast. The day follows by activities like studying, working or traveling. During the lunch, one should consume nutritious meal [...]. Dinner should consist of a light meal. Before going to bed, one should sit back and relax. By following the dincharya of Ayurveda, one can ensure a healthy life.

Though it is difficult to follow a stringent dinacharya in this fast moving life, it is highly recommended by Ayurvedic physicians, because a number of health benefits are associated with it. The dinacharya makes one to lead a healthy and disciplined life. According to the latest studies in the field of medical science, people who stick to the daily routine are more fit than those, who do not have a particular time to perform their everyday activities. It is said that dinacharya reduces the stress level to a great extent. In addition to this, the person's body is purified and detoxified. Therefore, barring a few exceptions like sickness, very old and young age, Ayurvedic dinacharya is recommended for everyone.

Extract from http://ayurveda.iloveindia.com/dincharya/index.html

Please recall a situation when you behaved similar to the description of the Dincharya [Daily Routine]. Describe the situation in a few sentences below:

Male Identity: Experiments 1, 2, 1.A, 1.B.

The Masculine Side

The Masculine Side deals with the strength of the self. It is what causes you to act either timidly or self-confidently. The thing that is most important in determining the strength of the masculine side, is the value that you, at a deep level, place on yourself. This is a value you know within yourself that you have really and truly earned. It could be thought of as a sort of self esteem. Placing a high value on yourself affects your whole being and helps you feel strong and confident in operating your life. And, in the reverse direction, when you are able to operate your life confidently, things can really turn around for you because you get more out of life, and this automatically makes you place a higher value on yourself. You can build the masculine side through progress and small wins, through positive reinforcement, by practicing, and by doing things and generally taking an active part in operating your life.

If you have a strong masculine side, you are in charge of your own life because you are internally controlled. You tend to look people in the eye. You stand straight, and you usually command attention when you walk into a room, whether you say anything or not. This happens because of the strength within. [...] if you have a strong masculine side you are self-confident, and don't feel it is necessary to show off.

The masculine side is full of things that you have to be strong and self-confident in order to do. These include being able to claim your basic rights, such as the right to feel free to operate independently of others, and the right to belong or fit into society in any way you please. Claiming your rights also includes being able to stand up to people who try to take away your rights, either by force or intimidation, or by manipulation, or by trying to hinder you in choosing your own direction in life. The masculine side also includes the ability to take risks when appropriate, to be decisive when necessary, and to focus or concentrate in order to get something done. In addition, it includes being able to figure out how to accomplish things so you can get more of what you want out of life. Part of this is being able to figure out how to operate your life in a responsible manner, how to reason without distorting reality and without fooling yourself, and how to accurately weigh probabilities so that you know the most likely outcome to expect in situations you come across.

Extract from ***http://www.lovesedona.com/02.htm.***

Please recall a situation in which you behaved according to the Masculine Side, as described in the reading.

You should try to recall that situation, your feelings and your thoughts in that occasion.
You should try to write a short essay (5-10 sentences), adding any details or considerations which come to your mind.

Female Identity: Experiments 1, 2, 1.A, 1.B.

The Feminine Side

[...] the Feminine Side is based, [...] on a value that you place on others. It could be thought of as a sort of other esteem. The value you place on others affects your whole being. If you have a strong feminine side and place a high value on others, you are often giving and unselfish. You usually know what is good for people, and you tend to operate in ways that help others get what they want out of life. You happily let people operate their own life without interference from you, but when asked, you are also willing to help by supporting, cooperating, and giving advice. People feel comfortable with you because you give them who you are without pushing yourself on others.
If you have a strong feminine side, people also feel comfortable being around you because there is no selfishness for them to detect. [...]

If you have a strong feminine side, you often behave in ways that are considered feminine in nature. You do things you have to be giving and unselfish in order to do. These include recognizing the basic right of all people to use their own will to operate their own life, for example, by allowing them freedom to operate independently, freedom to fit in where and how they want, and freedom to choose what things to confront or face up to in life. Allowing people their basic rights also includes allowing them to control their own life without interference from you, to choose their own obligations in life without being manipulated by you, and to choose their own path or direction in life without hindrance from you. The feminine side also includes having enthusiasm and zest for life, and recognizing what things are worth getting enthusiastic about. And it includes having the persistence and tenacity to stay with things to the end, while still knowing when to give up on something if your energy is better used elsewhere.

In addition, the feminine side also includes being kind, compassionate, patient, responsive to the needs of others, and it includes knowing how much energy you can put into each of these without hurting yourself by draining your own energy.

Extract from ***http://www.lovesedona.com/02.htm.***

Please recall a situation in which you behaved according to the Feminine Side, as described in the reading.

You should try to recall that situation, your feelings and your thoughts in that occasion.
You should try to write a short essay (5-10 sentences), adding any details or considerations which come to your mind.

Power Identity: Experiments 3, 1.A.

Please recall a particular situation in which you had POWER over another individual or individuals. By power, we mean a situation in which you controlled the ability of another person or persons to get something they wanted, or were in a position to evaluate those individuals.

You should try to recall the situation, your feelings and your thoughts in that occasion. Please describe this situation in which you had power, what happened, how you felt, your thoughts, etc. in as much detail as possible below. You should try to write a short essay (5-10 sentences), adding any details or considerations which come to your mind:

Lottery-Choice Task: Experiments 1, 3;

In the next five screens, you will be offered choices between two alternative lotteries, and you will have to pick the one you prefer to play.

Your bonus will entirely depend on the choices you make about the lotteries and on chance, so please think carefully about your decisions.

There is no right or wrong choice, nor any rule of thumb: you have to choose the lottery you want to play based on your own preferences. Please note than lotteries can only make you gain money, and never lose anything!

There are ten choices to be made for each screen. Here is an example of a choice:

20%, $2.5; ----- 80%, $1	○	○	20%, $3.3; ----- 80%, $0.1

The lottery on the left pays $2.5 with a probability of 20%, and $1 with a probability of 80%. The lottery on the right pays $3.3 with a probability of 20%, and $0.1 (10 cents) with a probability of 80%.

Think about it this way: there are five balls in an urn, numbered 1,2,3,4,5. You have to pick one number, and you choose 2. Then one ball is randomly extracted from the urn. If ball number 2 is extracted, the lottery to the left pays $2.5; if any other ball is extracted (1,3,4 or 5), it pays $1. The same applies to the lottery to the right.

For each line, you will have to fill the circle next to the lottery you prefer to play.
If you fill the left circle, you prefer to participate in the lottery to the left compared to the one to the right.
If you fill the right circle, you prefer to participate in the lottery to the right compared to the one to the left.
You will need to complete all the set of choices before moving to the next screen.

In each screen, lottery payoffs are the same, BUT the probabilities with which each payoff is paid change. They always change by the same amount at each row going down.

To determine your bonus, one of the five screens will be picked at random. Then, one of the ten choices in the screen will be picked at random. The lottery you have chosen will be played by generating a random number. You will earn the payoffs shown DIVIDED by one hundred.

Since the line which will determine your bonus is picked randomly AFTER you complete the whole task, no one knows which one it will be in advance. It is in your own best interest that you make each decision as if every lottery you choose will be played.

Each screen should take no more than 2 minutes to be completed.

○ I am ready to start with the lottery choices.

For each line, please pick one choice by filling the corresponding circle:

$6 for sure		50% -- $7
$5.7 for sure		50% -- $7
$5.4 for sure		50% -- $7
$5.1 for sure		50% -- $7
$4.8 for sure		50% -- $7
$4.5 for sure		50% -- $7
$4.2 for sure		50% -- $7
$3.9 for sure		50% -- $7
$3.6 for sure		50% -- $7
$3.3 for sure		50% -- $7
$3 for sure		50% -- $7
$2.7 for sure		50% -- $7
$2.4 for sure		50% -- $7
$2.1 for sure		50% -- $7
$1.8 for sure		50% -- $7
$1.5 for sure		50% -- $7
$1.2 for sure		50% -- $7
$0.9 for sure		50% -- $7
$0.6 for sure		50% -- $7
$0.3 for sure		50% -- $7
$0 for sure		50% -- $7

Investment Task:

Direct Investment: Experiment 1.A;

In each of the next 5 screens, you will be assigned a certain amount of money, you will be described an investment opportunity, and asked two questions. First, whether you want to invest any of the money you have been assigned. Second, if yes, how much you want to invest.

For each investment, there is a given probability that the investment gains and that it fails. If it gains, you

will keep the invested money and gain something. If it fails, you will lose the invested money. After you make your decisions, a computerized random process will tell you if the investment gained or failed. The process will say "gain" or "fail" based on the probabilities which are given to you, and it is the same as rolling dice (which we cannot do on the internet).

Your answers in this section will contribute to determine your bonus.
We will sum up all the money you have after each investment. When payment is processed, you will be given the final amount DIVIDED by ten thousands as part of your bonus. Hence, it is in your best interest to make all decisions carefully, given they all contribute to determine your payment.

○ I am now ready to move to the investment decisions.

You are now given $100.

If you invest any money, the investment will succeed with probability 1/2. For instance, if we roll a die, the investment succeeds when the die comes up 2, 4 or 6. It fails when the die comes up 1,3 or 5.

When the investment succeeds, you get back 3 times your investment. If it fails, you lose it. For instance, if you decided to invest $30 and the investment succeeds, you will be given back $90. In addition, you will keep the $70 you did not invest. Hence, at the end, if the investment succeeds you will have $160.

If the investment fails, you will lose the $30, and will be left with the $70 you have not invested.

If you do not invest anything, you will be left with the initial $100 in any case.

	Yes, I do want to invest some money	No, I do not want to invest at all
Do you want to invest any money?		

If yes, how much money do you want to invest? Please write the number below (please leave a blank if you decided not to invest). You can pick any amount below or equal to 100.

I want to invest $

Delegated Investment: Experiment 1.B.

Now imagine that Sally, a lady in your neighborhood, hires you to administer some of her money.

She would like to invest some money and gain something, instead of keeping it in her checking account.

In each of the next 6 screens, you will be given $100 by Sally.
At the same time, you will be made aware of an investment opportunity. You will be told the probability that the investment succeeds or fails. You will have to decide if you want to invest or not.

If you invest and the investment succeeds, Sally will get back more than the money she gave you.

If you invest and lose, Sally will get back whatever is left.

If you do not invest, Sally will get back the amount she gave you at the beginning, which to her is the same as having kept it in her checking account.

After you make your investment decision, a random number generator will tell you if the investment succeded or failed, based on the probabilities you are told. This is the same as rolling dice, which we cannot do on the internet.

Your answers in this section will determine your bonus - think of it as Sally's remuneration for your services.
We will sum up all the money you have after each investment. When payment is processed, you will be given the final amount DIVIDED by ten thousands as your bonus. Hence, it is in your best interest to make all decisions carefully, given they all contribute to determine your payment.

I am now ready to move to the investment decisions.

Online Appendix B: Additional Results

Figure B.1: **Age Distribution of Subjects across the Experiments**

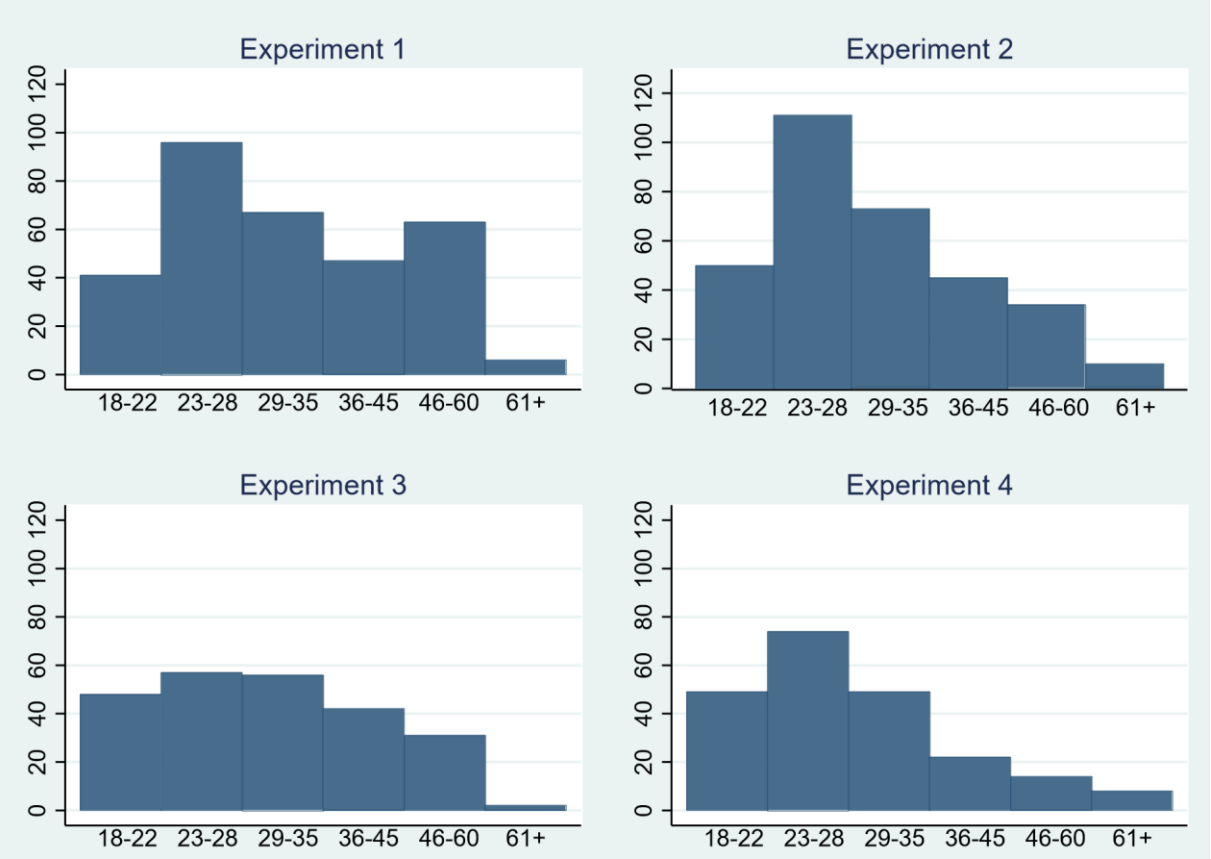
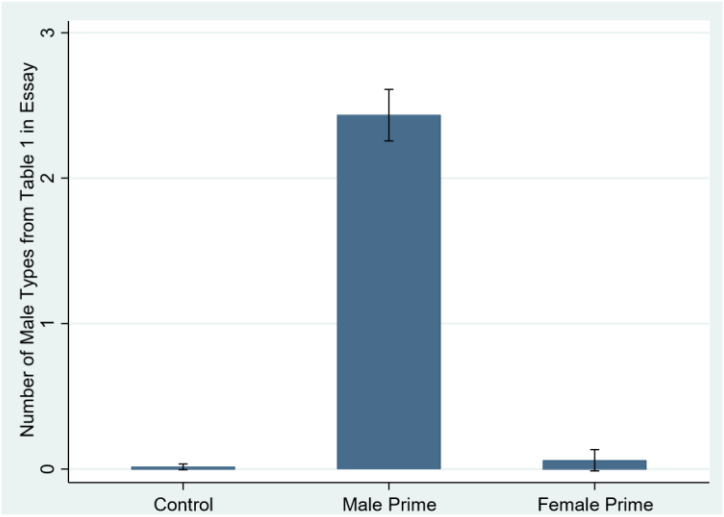


Figure B.2: **Manipulation Check for Experiment 1**

(a) Average Number of Male Types from Table 1 reported in Experiment 1



(b) Average Number of Female Types from Table 1 reported in Experiment 1

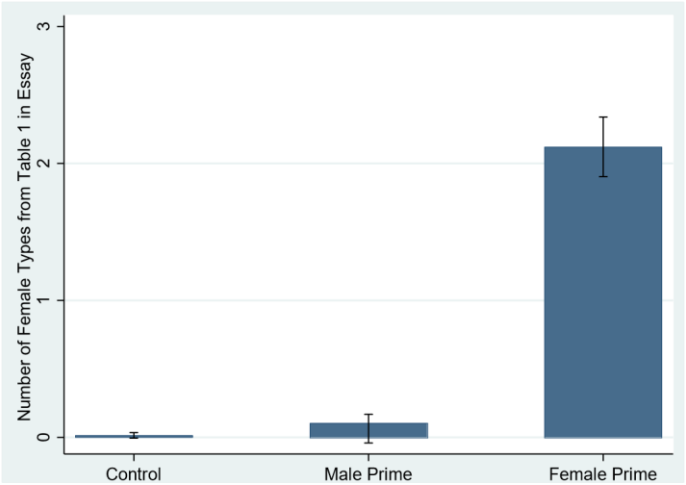


Figure B.2 reports the results for the manipulation check for Experiment 1. Panel (a) shows the average number of times each subject reported any of the types related to male individuals enlisted in Table 1 in their essays, by experimental conditions. Panel (b) shows the average number of times each subject reported any of the types related to female individuals enlisted in Table 1 in their essays, by experimental conditions. The bars represent 95% confidence intervals for the arithmetic mean within each group of subjects.

Figure B.3: **Risk Tolerance Elicitation**

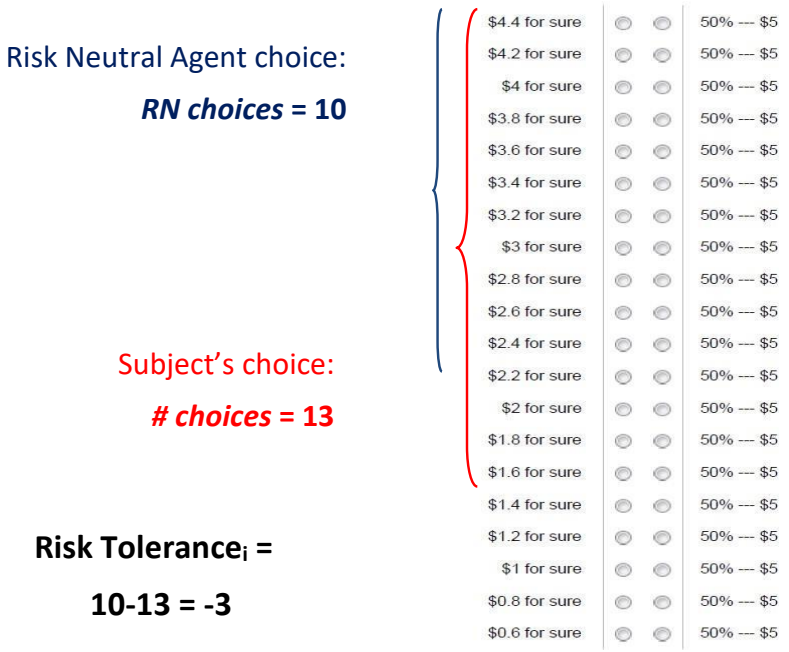
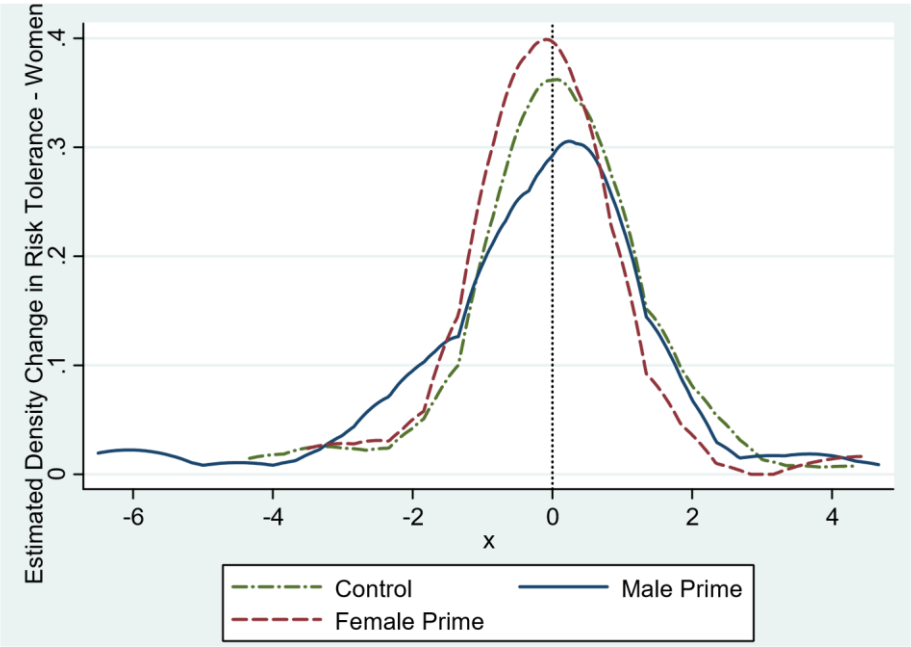


Figure B.3 describes the construction of the Risk Tolerance measure based on the subject's choices. Subjects face a set of screens, each including a variable number of choices between a decreasing certainty equivalent and a fixed lottery. The Risk Tolerance measure is the difference between the number of times a risk neutral agent would choose the certainty equivalent in the screen, and the number of times the subjects chooses the certainty equivalent in the screen. Note the outcome of interest is the *change* in the risk tolerance measures elicited before and after the experimental manipulations within subjects, and hence specificities of the same elicitation method used in both elicitation instances are not crucial to the results.

Figure B.4: **Risk Tolerance Change and Beliefs in the Female Subsamples**

Panel A depicts estimated densities for the change in risk tolerance measured at the subject level across experimental conditions in Experiment 1 (Identity prime) for female subjects. The change in risk tolerance is measured as follows: $\Delta RiskTolerance_i = RiskTolerance_{post,i} - RiskTolerance_{pre,i}$, where $RiskTolerance_{pre,i}$ and $RiskTolerance_{post,i}$ are elicited via lottery choices a la Holt and Laury (2002) before and after exposure to the experimental conditions, respectively. Panel B depicts the average measure of "better-than-average" beliefs constructed in Experiment 4 for the female subsample across experimental conditions.



A. Distribution of ΔRisk Tolerance (female subsample)

B. Better-than-Average Beliefs (female subsample)

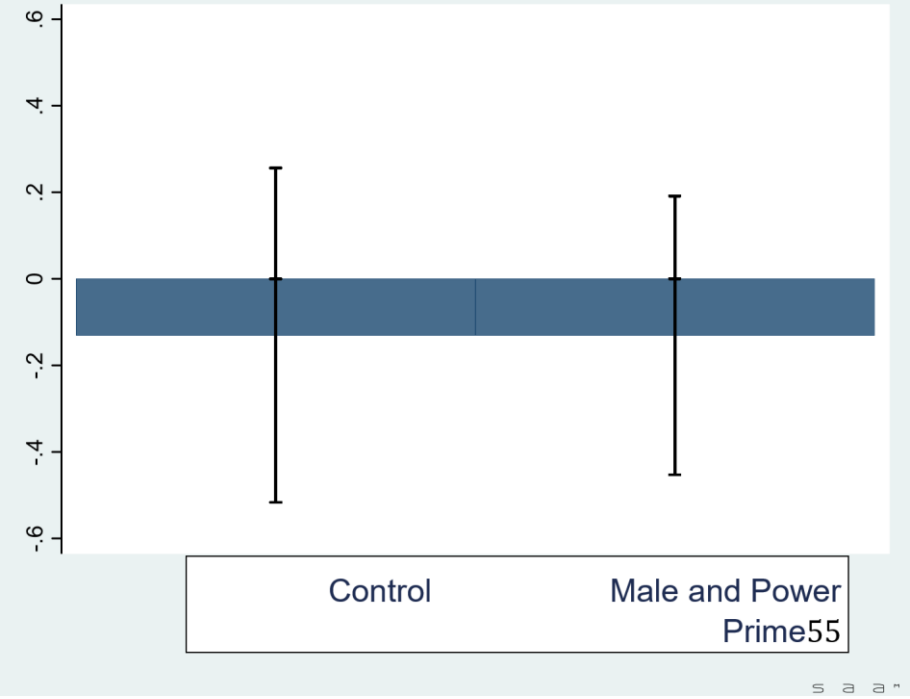


Figure B.5: **Identity and Overconfidence cues in financial products ads**

The set of pictures below are examples of financial products ads across countries and for different types of products with male identity and overconfidence cues.



Table B.1: **Experiment 1 - Contribution of right tail and average to results**

Panel A shows results for the change in risk tolerance in the male subsample in Experiment 1. Panel B shows results for the change in risk tolerance for the full sample of subjects in Experiment 1. The outcome variable (Δ*RiskTolerance*) is winsorized at the 5-95 percentiles in the left Panels, at the 10-90 percentiles in the right panels. Hubert-White s.e. are reported in brackets. Statistical significance is marked as follows: ** 5%, * 10%.

A. Experiment 1: Winsorizing outliers (male subsample)

Male Only	(1)	(2)	(3)	(4)	(5)	(6)
		Winsorize 5-95 perc.			Winsorize 10-90 perc.	
M. Prime	0.319	0.336	0.335	0.293	0.306	0.304
	(0.197)	(0.195)*	(0.195)*	(0.165)*	(0.161)*	(0.163)*
F. Prime	0.481	0.447	0.450	0.419	0.377	0.373
	(0.208)**	(0.215)**	(0.221)**	(0.189)**	(0.194)*	(0.199)*
Mean omitted		-0.170			-0.133	
Median omitted		0			0	
Session f.e. Age group f.e.	X	X		X	X	
Education f.e.			X			X
Observations	148	148	X		148	X
R_2	0.03	0.09	148	148	0.09	148
			0.09	0.04		0.09

B. Experiment 1: Winsorizing outliers (full sample)

Full Sample	(1)	(2)	(3)	(4)	(5)	(6)
		Winsorize 5-95 perc.			Winsorize 10-90 perc	
M. Prime	-0.265	-0.269	-0.237	-0.189	-0.188	-0.169
	(0.194)	(0.196)	(0.196)	(0.154)	(0.155)	(0.156)
F. Prime	-0.117	-0.162	-0.153	-0.088	-0.127	-0.122
	(0.202)	(0.209)	(0.209)	(0.162)	(0.167)	(0.167)
Male*M. Prime	0.590	0.588	0.566	0.489	0.484	0.469
	(0.276)**	(0.277)**	(0.278)**	(0.225)**	(0.225)**	(0.227)**
Male*F. Prime	0.628	0.650	0.643	0.543	0.557	0.553
	(0.276)**	(0.280)**	(0.281)**	(0.235)**	(0.239)**	(0.241)**
Male	-0.253	-0.262	-0.271	-0.244	-0.248	-0.252
	(0.179)	(0.179)	(0.182)	(0.151)	(0.151)	(0.154)
Mean omitted		0.050			0.072	
Median omitted		0			0	
Session f.e. Age group f.e.	X	X		X	X	
Education f.e.			X			X
Observations		320	X		320	X
R_2	320	0.04	320	320	0.05	320
	0.03		0.06	0.03		0.10

Table B.2: **Investment Decisions by the Female Subsamples**

Panel A shows results for the investment decisions made by women in Experiments 1.A and 1.B. The left panel reports marginal effects computed from the following probit model:

$$Pr(Invest = 1)_{iae} = \Phi(\alpha + \beta \times MalePrime + \gamma \times OverconfidencePrime + \eta_a + \eta_e)$$

where *Invest* is 1 if subject *i* in age group *a* and education level group *e* invests any positive amount, zero otherwise. Φ(.) is the normal cdf. *MalePrime* and *OverconfidencePrime* equal one for subjects exposed to the male identity or power prime; η_a and η_e are fixed effects for subjects' age and education level groups for female subjects. The right panel reports results for a tobit model whose dependent variable is the amount of money a subject invests in each opportunity, which is censored at $0 and $100. Standard errors are clustered at the subject level, and computed with the delta method. Statistical significance is marked as follows: *** 1%, ** 5%, * 10%.

A. Experiment 1.A: Investment Decisions (female subsample)

Female subsample	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Probability of Investing				Amount Invested			
Male Prime	0.113	0.137	0.135	0.136	3.554	5.638	5.961	5.763
	(0.076)	(0.074)*	(0.074)*	(0.073)*	(5.667)	(5.321)	(5.384)	(5.359)
Overconfidence Prime	-0.026	-0.017	-0.004	-0.003	-7.751	-6.813	-6.616	-6.778
	(0.073)	(0.073)	(0.074)	(0.073)	(6.060)	(5.851)	(5.864)	(5.848)
Constant					11.94	14.87	-13.54	2.803
					(4.48)***	(5.64)***	(15.82)	(15.32)
Age group f.e.		X	X	X		X	X	X
Education f.e.			X	X			X	X
Success prob. f.e				X				X
Observations	387	387	384	384	387	387	384	384
(pseudo-)R^2	0.01	0.02	0.02	0.09	0.01	0.01	0.01	0.03

B. Experiment 1.B: Investment Decisions (female subsample)

Female Subsample	(1)	(2)	(3)	(4)	(5)	(6)
	Probability of Investing			Amount Invested		
Male Prime	0.005	0.001	0.004	0.54	0.50	0.84
	(0.055)	(0.053)	(0.052)	(5.37)	(5.22)	(2.18)
Constant				38.26		
				(3.92)***		
Age group f.e.		X	X		X	X
Education f.e.			X			X
Observations	230	230	228	230	230	228
(pseudo-)R^2	0.00	0.03	0.06	0.00	0.01	0.01

Table B.3: **Just Post It - Evidence vs. Fabricated Data**

Panel A refers to the full sample, Panel B to the male subsample. For each cell the sample mean is reported, the standard deviation is reported in round parentheses, and the number of subjects in curly parentheses. Panel C shows results for running the test of Simonsohn (2013) for each experiment and sample. The test consist of: (i) estimating 100,000 samples randomly drawn from a normal distribution with mean equal to the one in the condition, and standard deviation equal to the average across conditions in each experiment (samples are drawn from a binomial distribution with probability equal to the mean of the extensive margins in Experiment 1.A and Experiment 1.B); (ii) Computing the standard deviation of the standard deviations in each drawn sample; (iii) reporting the ratio of times that the simulated standard deviation of standard deviations are lower than the actual one in the data.

A. Summary Statistics (full samples)

	Male Identity Prime	Female Identity Prime	Power Prime	Success Prime	Control
Experiment 1	-0.01 (1.74) {125}	0.05 (1.07) {66}			-0.06 (1.22) {129}
Experiment 3			-0.21 (1.43) {107}	-0.38 (1.37) {109}	-0.51 (1.62) {107}
Experiment 1.A (extensive)	0.747 (0.436) {237}		0.680 (0.468) {231}		0.624 (0.485) {234}
Experiment 1.A (intensive)	27.95 (29.81) {237}		23.34 (28.55) {231}		23.41 (29.40) {234}
Experiment 1.B (extensive)	0.904 (0.296) {208}				0.852 (0.356) {216}
Experiment 1.B (intensive)	49.93 (32.17) {208}				42.04 (30.07) {216}
Experiment 2	0.079 (1.711) {214}				-0.090 (1.743) {110}

B. Summary Statistics (male subsamples)

	Male Identity Prime	Female Identity Prime	Overconfidence Prime	Success Prime	Control
Experiment 1	0.30 (1.54) {61}	0.22 (0.81) {32}			-0.21 (1.09) {55}
Experiment 3			0.03 (1.32) {54}	-0.47 (1.39) {59}	-0.49 (1.40) {56}
Experiment 1.A (extensive)	0.781 (0.416) {114}		0.823 (0.384) {96}		0.648 (0.480) {105}
Experiment 1.A (intensive)	34.92 (35.42) {114}		35.47 (33.97) {96}		26.99 (33.41) {105}
Experiment 1.B (extensive)	0.949 (0.221) {98}				0.844 (0.365) {96}
Experiment 1.B (intensive)	60.88 (35.03) {98}				44.96 (32.75) {96}
Experiment 2	0.405 (1.407) {84}				-0.024 (1.968) {41}

C. P-values for random draw of samples across conditions

	Full Sample	Male Subsample
Experiment 1	0.99991	0.99982
Experiment 3	0.82831	0.10523
Experiment 1.A (ext. margin)	0.45736	0.41635
Experiment 1.A (int margin)	0.20449	0.17365
Experiment 1.B (ext. margin)	0.50212	0.47808
Experiment 1.B (int. margin)	0.67175	0.49040
Experiment 2	0.17617	0.98543