Is it lie aversion, risk-aversion or tax audit aversion? Modeling deception under risk and no risk

Tei Laine (tlaine@msh-alpes.fr)
Maison des Sciences de l'Homme-Alpes, Université Pierre Mendès-France
1221 avenue centrale, Domaine universitaire, St Martin d'Hères, France

Tomi Silander (tomi.silander@xrce.xerox.com)
Xerox Research Centre Europe, 6 chemin de Maupertuis
38240 Meylan, France

Kayo Sakamoto (sakamotok@ihpc.a-star.edu.sg) **Ilya Farber (farberi@ihpc.a-star.edu.sg)**
Institute of High Performance Computing, A*STAR, 1 Fusionopolis Way, #16-16 Connexis North, Singapore 138632

Abstract

We studied deceptive decision making in hypothetical scenarios that involved risk of being caught of deceiving, or a penalty after being caught of deceiving, or both. We found that the deception rate was the lowest in the scenarios involving both the risk and the penalty. Our hierarchical model for deception suggests that in balancing the possible benefits from deception, the personal discomfort of getting caught is as large or larger than the inherent aversion to deception.

Keywords: Decision making; risk attitudes; deception; incentives; MTurk.

Introduction

In our recent study that asked participants' choices between risky and certain options (in the context of tax return), which either involved deception (deception condition) or did not involve deception (gamble condition), we observed a high rate of deception aversion[1] in the condition in which deception resulted in the better outcome than being honest, but involved both risk, i.e., non-zero probability of detection, and a potential loss in the form of a tax penalty, if the deception was detected (Laine, Sakamoto, & Silander, 2013). The participants who were particularly deception averse in the deception condition were also more risk averse than others in the gamble condition, which had equivalent risky and certain outcomes, but involved no deception.[2] In other words, the participants who took more risk in gamble condition, also deceived more in the deception condition.

We speculated if the reason for such a high level of tax compliance in the risky deception condition was exceptionally high level of risk aversion or exceptionally high level of deception aversion, or alternatively the task domain combined with the participant pool characteristics. We used Amazon MTurk workers from the US. This is a group of individuals who are willing to do simple tasks for little monetary compensation. Alternatively, based on their own prior experiences or knowledge of others' encounters with the Internal Revenue Service (IRS), our participants (most of them US tax payers) may have wanted to avoid any friction (even hypothetical) with the tax authorities, and indicated their willingness to pay due taxes, even in the presence of substantial financial incentives for evasion.

To rule out the explanation pertaining to the task domain and the participant pool, we conducted another study with MTurk participants, and added conditions from which we excluded either the risk or the tax penalty. Again we observed an exceptionally high rate of deception aversion in the condition that involved both a non-zero probability of detection and and a penalty after the deception was detected. Thus, it seems that the "IRS aversion", in other words the aversion to a potential audit by the tax authorities, is not alone enough to explain the high level of tax compliance, since the participants demonstrated some willingness to evade taxes in conditions from which either the risk or the penalty for detected deception was absent. In this study we wanted to find out what distinguishes those participants who refused to deceive in their taxes no matter what from those who properly incentivized switched from complete tax compliance to some degree of tax evasion.

Who are those who do X, where X $\in$ {take risk, deceive, evade taxes}

In general, people tend to be risk averse when facing gains and risk seeking when facing losses (Holt & Laury, 2002; Kahneman & Tversky, 1979). Many economic models of choice behavior are based on the concept of individual risk attitude, which can be measured experimentally and modelled with the shape and parameters of a utility function (Holt & Laury, 2002; Isaac & James, 2000; Weber, 1998). It has been considered a stable construct similar to personality traits, which drives behavioral patterns across situations (Blais & Weber, 2006).

However, this interpretation is problematic, since several studies have found that the risk attitude varies across task types (e.g., hypothetical vs. real outcomes) (Holt & Laury, 2002; Taylor, 2013), domains (e.g., financial or health related decision making), elicitation methods (e.g., choice between

[1] In this experiment 279 (42%) out of 672 participants did not choose the risky deceptive option a single time in the deception condition.

[2] Only 25 out 672 participants never chose the risky option in the gamble condition.

gambles, a questionnaire, an auction, or a multiple price list method (Berg, Dickhaut, & McCabe, 2005; Charness, Gneezy, & Imas, 2013; Crosetto & Filippin, 2013)), and even ages (Harbaugh, Krause, & Vesterlund, 2002). For instance, Isaac and James (2000) demonstrated in a within participant study using two different risk elicitation tasks that some participants could in an instant turn from being risk averse to risk seeking, whereas some others remained risk neutral in both tasks. Blais and Weber (2006) suggest that even if an individual's risk attitude towards a perceived risk does not differ from one domain to another or from one situation to another, her risk-taking behavior might, if she perceives the risk and the benefits to be different in those situations.

There are differences between genders, too. Harris, Jenkins, and Glaser (2006) found that women's lower engagement rate in risky activities correlates with their tendency to judge negative events more likely and expected enjoyment of risky activities less highly than men in gambling, recreation, and health domains. However, there was no gender differences in risk-taking and risk perception in the social domain. Mixed findings have been reported on age differences in risk-taking propensity. Many studies have found more risk aversion in older people, but also the opposite in certain circumstances, or even no differences between age groups (see for instance: Deakin, Aitken, Robbins, & Sahakian, 2004; Dror, Katona, & Mungkur, 1998; Harbaugh et al., 2002; Huang, Wood, Berger, & Hanoch, 2013; Mather, 2006).

How do these studies relate to age and gender differences in deceptive behavior? Studies with adults have shown that while women are more lie averse than men in general, they are more likely to lie if it benefits others, and less likely if it hurts others, whereas men lie more for self-serving purposes, particularly monetary gains (dePaulo, Kashy, Kirkendol, & Wyer, 1996; Dreber & Johannesson, 2008; Erat & Gneezy, 2012). Finally, even if there is considerable heterogeneity in tax evasion within any group defined by a demographic category such as income or age, studies have shown that men evade taxes more than women, high income people evade taxes less than low income people, and married and under 65-year-old tax payers evade more than others (Slemrod, 2007).

These findings at least partially suggest that there is a link between risk-attitude and propensity to deceive, even to certain extent in taxes. However, most studies have focused on individuals or groups who partake in behaviors or activities in question (acts of commission), whereas in the following we are interested in those who do not (acts of omission).

Experimental design

According to the standard economic model of rational and selfish human behavior (i.e., the "homo economicus"), one should deceive if it is beneficial compared to being honest, and the decision should be solely determined by the trade-off between the gain from lying and the cost incurred if detected, given the probability of detection (Abeler, Becker, & Falk, 2014; Gneezy, Rockenbach, & Serra-Garcia, 2013). Therefore, the policies to curb deception, for instance in tax returns, have almost solely focused on increasing the detection probability and the penalty. However, these measures are not necessarily effective if people's deceptive behavior is driven by internal rewards instead of cost-benefit analysis of external rewards (Mazar, Amir, & Ariely, 2008); sometimes otherwise inconceivable indisposition to deceive have been attributed to factors like pure lie aversion (Fischbacher & Heusi, 2008; Gneezy, 2005; Gneezy et al., 2013; Erat & Gneezy, 2012; López-Pérez & Spiegelman, 2012; Lundquist, Ellingsen, Gribbe, & Johannesson, 2009), altruism (Abeler et al., 2014), maintenance of positive self-image, e.g., avoiding to appear greedy (Mazar et al., 2008; Fischbacher & Heusi, 2008), and moral considerations based on the norms and values of the society (Mazar et al., 2008; Sip et al., 2012).

Since our primary goal was to study the role of risk and monetary incentives in deception, we wanted to rule out the above factors. Instead, to more efficiently isolate the effect of risk from the effect of the outcomes, we made the expected value of the risky deceptive option higher than the expected value of the non-deceptive option, and added two conditions in which—still maintaining the expected value difference—either (1) both deception and being honest resulted in a certain outcome, i.e., there was no risk, or (2) failing the deception resulted in the same outcome as being honest, i.e., there was no penalty for detected deception.

Method

Participants We recruited 372 participants in Amazon MTurk to complete an online questionnaire in Qualtrics. After discarding data from participants who either did not complete the whole questionnaire or failed the attention check question we had 301 participants (130 women, 170 men, one unknown; median age 30 years, age range: 18-71 years). All participants were native English speakers, aged 18 or above, residing in the US. Each participant received USD 1.00 for their participation.

Material Using filling in a tax return as the deceptive domain, we prepared six gain-facing and six loss-facing scenarios in three conditions:

VarRisk Variable risk and constant outcome from successful deception (gain of $1000 and loss of $0), and from being truthful or unsuccessful deception (gain of $0 and loss of -$1000). We used detection probabilities of 0.01, 0.1, 0.25, 0.5, 0.75, and 0.90.

VarOut Variable outcomes and constant risk at 0% detection probability.

VarBoth Both variable outcomes and variable risk.

The **VarBoth** condition matched the **VarOut** condition in the expected values of options (see Table 1), and the **VarRisk** condition in the detection probabilities. This design allowed us to address the effects of risk and expected utility separately.

Example scenarios in **VarBoth** and **VarOut** in the gain condition are shown in Figure 1.

Table 1: Expected values in **VarBoth** and **VarOut** conditions for gains and losses.

Gains	Honest option $	891	810	675	450	225	90
	Deceptive option $	990	900	750	500	250	100
Losses	Honest option $	11	110	275	550	825	990
	Deceptive option $	10	100	250	500	750	900

VarBoth scenario Jenny is doing her yearly taxes. In addition to her regular job, she has done some part-time work on the side. Jenny is aware that she needs to report all of her income in her tax return.

If she omits the extra part-time earnings from her report, she will get back $1,000. If she includes these extra earnings in her report, she will get back $675. The part-time work was done for a small business, which might or might not report this payment to the IRS. If Jenny fails to declare the income, but the business reported it, she will lose her entire tax refund as a penalty, meaning that she will get $0. But since there's no way to prove that it wasn't just an accounting mistake, there is no other fine or long-term penalty for getting caught.

From talking to other people who've worked at the same place, Jenny estimates that there is a 25% chance that the business reported the payment. If you were Jenny, would you

[] Report the extra income and get a refund of $675.

[] Not report the extra income, in which case you have a 75% chance of getting back $1,000, but also a 25% chance of getting back $0.

VarOut scenario John is doing his yearly taxes. In addition to his regular job, he has done some part-time work on the side. John is aware that he needs to report all of his income in his tax return. If he omits the extra part-time earnings from his report, he will get back $750. If he includes these extra earnings in his report, he will get back $675.

The part-time work was paid in cash and the employer didn't record his name. Therefore, John knows that there is no chance that the IRS knows about this income. If he leaves it out of the report, there is no chance that he will be caught. If you were John, would you

[] Report the extra income and get a refund of $675.

[] Not report the extra income and get a refund of $750.

Figure 1: Example questions in **VarOut** gain and **VarBoth** gain conditions.

Procedure After giving their informed consent the participants were asked to make their choices in six sets of six questions (the order of the sets was randomized for each participant.[3] All participants answered all 36 questions, so the experimental condition manipulation (**VarBoth** vs. **VarOut** vs. **VarRisk**) was within participant. After finishing the choice questionnaire they answer a set of 30 risk- and deception attitude questions and completed a brief numeracy test, results of which are not reported here. They finished by filling in optional background information, including age, gender, and education. The questionnaire ended with a debriefing. It took them 15 minutes on average to finish the whole experiment.

[3]We prepared two versions of each scenario, one with a female and one with a male tax payer, and picked one randomly for each participant.

Results

In both **VarOut** and **VarRisk** conditions we observed much more deception than in the **VarBoth** condition, see Figures 2 and 3.

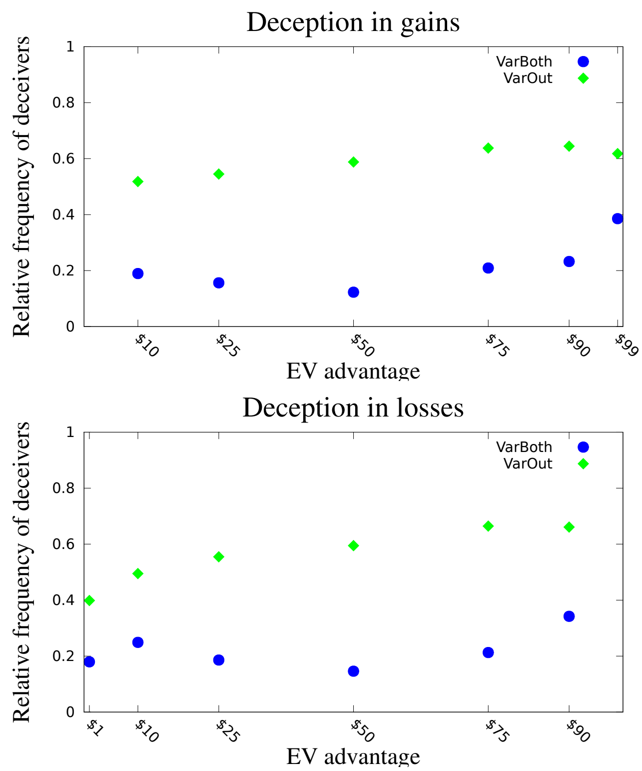

Figure 2: Relative frequencies of deceivers in the conditions varying only outcome, or both risk and outcome.

Risk aversion vs. penalty aversion The percentage of participants who never deceived was much higher in the **VarBoth** condition than in the other two conditions (see Table 2), and roughly corresponded to the percentage observed in Laine et al. (2013) (37% compared to 42%). The number of non-deceivers in the **VarBoth** condition significantly differed from the other two conditions both overall, $\chi^2(2, N = 301) = 20.0076$, $p = 4.523e{-}05$, and also separately for gains, $\chi^2(2,301) = 38.0499$, $p = 5.465e{-}09$, and losses, $\chi^2(2, N = 301) = 31.259$, $p = 1.63e{-}07$. All other differences were insignificant, i.e., between **VarRisk** and **VarOut** conditions overall ($p = 0.77$), for gains ($p = 0.65$), and for losses ($p = 0.64$).

The low deception rate in **VarBoth** condition suggests that

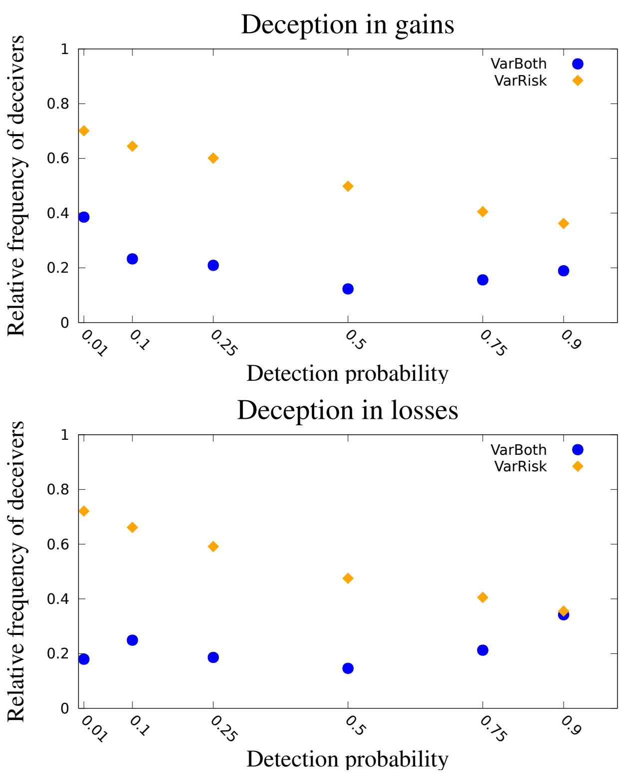

Figure 3: Relative frequencies of deceivers in the conditions varying only risk, or both risk and outcome.

Table 2: Number of deceivers and non-deceivers in each condition

Condition	Non-deceivers	Deceivers
VarOut	68	233
VarRisk	71	230
VarBoth	112	189

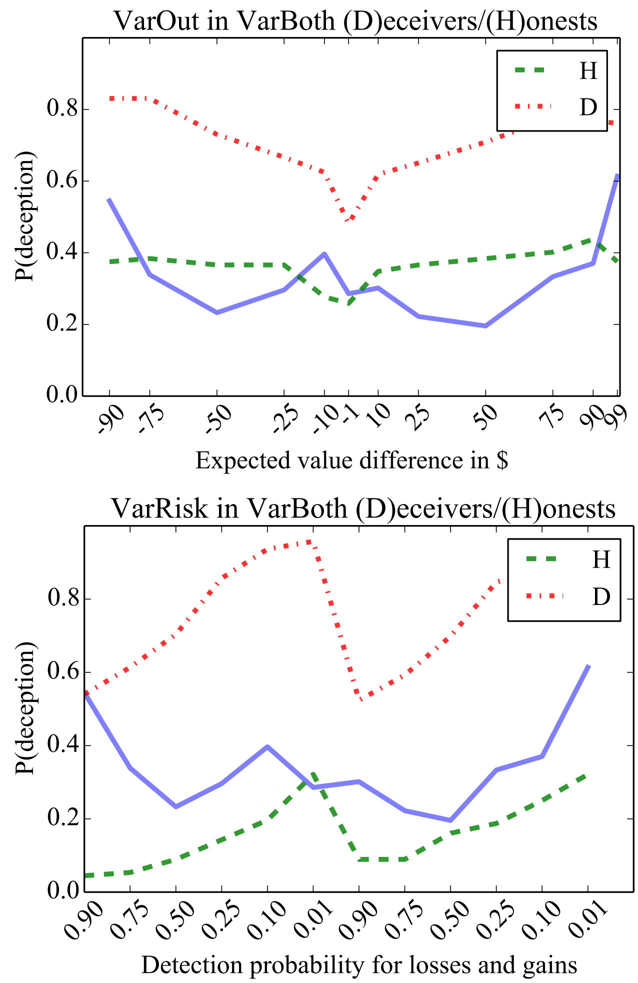

Figure 4: Probability of Deception in **VarOut** and **VarRisk** conditions for those who did and did not deceive in **VarBoth** condition.

our participants were extremely risk averse. This is partially supported by the higher deception rate observed in the riskless **VarOut** condition. On the other hand, the expected value differences between deceptive and non-deceptive options in these two conditions were equivalent, so one would expect the same (deceptive) choices in both of them, if the participants were basing their decisions on the expected values. It is less clear, though, why there is no significant difference between **VarOut** and **VarRisk** conditions, since the latter did involve risk.

Looking at the behavior in the other two conditions, we try to gain insight on what distinguishes those participants who never deceived in **VarBoth** condition (non-deceivers) from those who deceived at least once (deceivers). First, we did not find any meaningful patterns or relationships between the choices and the background variables (e.g., age or gender) in either group. Second, in both **VarOut** and **VarRisk** conditions deception rate was higher in **VarBoth** condition in both of these groups. However, there were some qualitative differences both between the groups and between the two conditions.

While removing risk (**VarOut** condition) increased the probability of deception equally in deceivers and non-deceivers, removing the penalty (**VarRisk** condition) increased it much more in deceivers (see Figure 4: the blue solid line represents the deceivers in **VarBoth** condition). In other words, incurring no potential penalty after the deception is detected seemed to be more effective incentive to deceive than having no risk of getting detected in the first place, but only for those who had a higher propensity to deceive to start with.

In both of these conditions the deceivers responded better to incentives; their probability of deceiving appears to be a function of detection probability in **VarRisk** condition and the expected value difference in **VarOut** condition, whereas such a pattern was not as apparent in non-deceivers. One can hypothesize that this is because the non-deceivers have a higher cost of deception which overrides the effect of incentives no matter how attractive they are.

A hierarchical model

These findings motivated us to entertain a Bayesian hierarchical model (Lee, 2011) where in addition to the expected monetary gain from deception each individual i has two factors that influence her choices. The first one is her own "prize" $cdec_i$, which measures the monetary equivalent of inherent cost of lying (i.e., pure deception aversion). The other one is the cost of getting caught, $cdet_i$, which measures the shame or regret of getting caught. The model assumes that these two factors vary in the population.

The data we use in this model consists of the responses of the 301 participants in six gain questions in **VarOut** condition and six gain questions in **VarBoth** condition. We denote the choice of the participant i in the question k, ($k \in 1,2,\ldots,12$) as D_{ik}, and the expected monetary value of the deception in question k as V_k. The probability of getting caught in question k is denoted as $pdet_k$. This probability is zero for the questions in **VarOut** condition. Finally, to turn the utilities (costs and expected gains) measured in terms of money into probabilities for choosing deception, we need a "temperature" parameter w that controls the mapping.

We can now express the model more formally,

$$
\begin{aligned}
\mu_{cdec} &\sim N(\mu_0, \sigma_0^2),\\
\sigma_{cdec} &\sim Uni(0, 1000),\\
cdec_i &\sim N(\mu_{cdec}, \sigma_{cdec}^2),\\
\\
\mu_{cdet} &\sim N(\mu_0, \sigma_0^2),\\
\sigma_{cdet} &\sim Uni(0, 1000),\\
cdet_i &\sim N(\mu_{cdet}, \sigma_{cdet}^2),\\
\\
w &\sim Uni(0, 10),\\
p_{ik} &= logit^{-1}(w \times (V_k + cdec_i + pdet_k \times cdet_i)),\\
D_{ik} &\sim Bernoulli(p_{ik}).
\end{aligned}
$$

The hyperparameters were set to non-informative values, and the estimation was conducted using PyMC python library that implements adaptive Metropolis sampling (Patil, Huard, & Fonnesbeck, 2010).

The posterior mean values for participants' detection and deception costs appear to indicate that on average they both play an equal role in deception (see Figure 5). The pure deception aversion seems to vary between \$30 and \$150, and the detection cost between \$40 and \$110. The joint density plot reveals that the individual's detection and deception costs tend to correlate, but there are people whose detection cost is twice their deception cost (say \$100 vs. \$50).

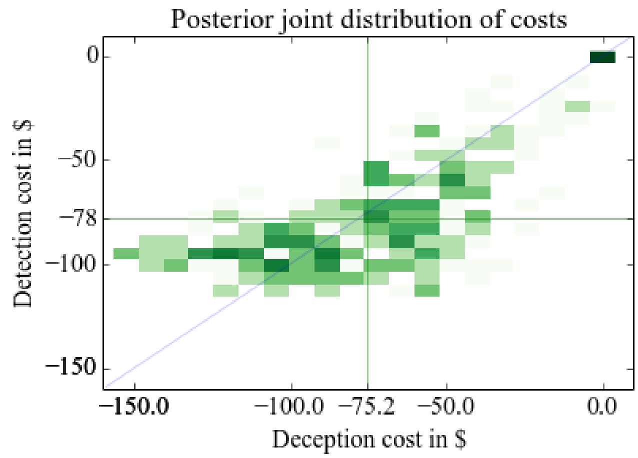

Figure 5: Joint posterior density of estimated deception and detection costs.

Discussion

It might be suggested that the MTurk workers, who are willing to spend time completing simple tasks to earn just few cents, must be unusual, and unusual in ways that can skew the effects of the experimental manipulations. They have been found to be more risk averse than other participant pools, for instance general public or student samples, but show the same pattern of risk attitudes by being risk seeking when facing losses and risk averse when facing gains (Horton, Rand, & Zeckhauser, 2011; Paolacci, Chandler, & Ipeirotis, 2010; Paolacci & Chandler, 2014; Rand, 2011).

Our participants also exhibited a seemingly high level of risk aversion. However, it was not the presence of risk per se that made the participants avoid deception, but also what may happen after their deception gets detected. If there was no difference in the outcomes when getting detected and telling the truth, the participants were willing to take risk and deceive, but if there was also a chance of incurring a penalty after getting detected, they were not. In other words, it was loss aversion rather than pure deception aversion that determined the deceptive behavior.

There are two possible interpretations of the findings: First, the participants who chose to deceive really were responding to monetary incentives, either to potentially higher gain when no risk was involved, or to the absence of loss (if getting detected) when risk was involved, rather than the presence of risk itself. However, the incentives seemed to influence more those participants who were already willing to deceive to some extent in the condition that involved both the risk and the penalty.

An alternative interpretation suggests that in those participants who never deceived both the pure lying cost and the cost of getting caught actually reflected the violation of social norms or individuals' own moral standards. In this case, the reluctance to deceive can be interpreted as maintenance of self-concept as suggested by Mazar et al. (2008).

References

Abeler, J., Becker, A., & Falk, A. (2014). Representative evidence on lying costs. *Journal of Public Economics*, *113*, 96–104.

Berg, J., Dickhaut, J., & McCabe, K. (2005). Risk preference instability across institutions: A dilemma. *Proceedings of the National Academy of Sciences*, *102*(11), 4209–4214.

Blais, A.-R., & Weber, E. U. (2006). A domain-specific risk-taking (dospert) scale for adult populations. *Judgment and Decision Making*, *1*(1), 33–47.

Charness, G., Gneezy, U., & Imas, A. (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization*, *87*, 43–51.

Crosetto, P., & Filippin, A. (2013). The "bomb" risk elicitation task. *Journal of Risk and Uncertainty*, *47*, 31–65.

Deakin, J., Aitken, M., Robbins, T., & Sahakian, B. J. (2004). Risk taking during decision-making in normal volunteers: Changes with age. *Journal of the International Neuropsychological Society*, *10*(4), 590–598.

dePaulo, B. M., Kashy, D. A., Kirkendol, S. E., & Wyer, M. M. (1996). Lying in everyday life. *Journal of Personality and Social psychology*, *70*(5), 979–995.

Dreber, A., & Johannesson, M. (2008). Gender differences in deception. *Economics Letters*, *99*, 197–199.

Dror, I., Katona, M., & Mungkur, K. (1998). Age differences in decision making: to take a risk or not? *Gerontology*, *44*(2), 67-71.

Erat, S., & Gneezy, U. (2012). White lies. *Management Science*, *58*, 723–733.

Fischbacher, U., & Heusi, F. (2008). Lies in disguise: An experimental study on cheating. *Journal of the European Economic Association*, *11*(3), 525–547.

Gneezy, U. (2005). Deception: The role of consequences. *The American Economic Review*, *95*(1), 384–394.

Gneezy, U., Rockenbach, B., & Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior & Organization*, *93*, 293–300.

Harbaugh, W. T., Krause, K., & Vesterlund, L. (2002). Risk attitudes of children and adults: Choices over small and large probability gains and losses. *Experimental Economics*, *5*, 53–84.

Harris, C. R., Jenkins, M., & Glaser, D. (2006). Gender differences in risk assessment: Why do women take fewer risks than men? *Judgment and Decision Making*, *1*(1), 48–63.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644–55.

Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: conducting experiments in a real labor market. *Experimental Economics*, *14*, 399–425.

Huang, Y., Wood, S., Berger, D., & Hanoch, Y. (2013). Risky choice in younger versus older adults: Affective context matters. *Judgment and Decision Making*, *8*(2), 179–187.

Isaac, R. M., & James, D. (2000). Just who are you calling risk averse? *Journal of Risk and Uncertainty*, *20*(2), 177–187.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

Laine, T., Sakamoto, K., & Silander, T. (2013). Do risk-averse people lie less? A comparison of risk-taking behavior in deceptive and non-deceptive scenarios. In M. Knauff, N. Sebanz, M. Pauen, & I. Wachsmuth (Eds.), *Proceedings of the 34th annual meeting of the cognitive science society* (pp. 2808–2813). Austin, TX.

Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, *55*(1), 1–7.

López-Pérez, R., & Spiegelman, E. (2012). Why do people tell the truth? Experimental evidence for pure lie aversion. *Experimental Economics*, *16*(3), 233–247.

Lundquist, T., Ellingsen, T., Gribbe, E., & Johannesson, M. (2009). The aversion to lying. *Journal of Economic Behavior & Organization*, *70*, 81–92.

Mather, M. (2006). A review of decision making processes: Weighing the risks and benefits of aging. In L. C. . C. Hartel (Ed.), *When I'm 64* (pp. 145–173). Washington, DC: National Academies Press.

Mazar, N., Amir, O., & Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, *45*(6), 633–644.

Paolacci, G., & Chandler, J. (2014). Inside the Turk: Understanding Mechanical Turk as a participant pool. *Current Directions in Psychological Science*, *23*, 184–188.

Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, *5*(5), 411–419.

Patil, A., Huard, D., & Fonnesbeck, C. J. (2010). Pymc: Bayesian stochastic modelling in python. *Journal of Statistical Software*, *35*(4), 1–81.

Rand, D. G. (2011). The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, *299*, 172–179.

Sip, K. E., Skewes, J. C., Marchant, J. L., McGregor, W. B., Roepstorff, A., & Frith, C. D. (2012). What if I get busted? Deception, choice, and decision-making in social interaction. *Frontiers in Neuroscience*, *6*, 1–10.

Slemrod, J. (2007). Cheating ourselves: The economics of tax evasion. *Journal of Economic Perspectives*, *21*(1), 25–48.

Taylor, M. P. (2013). Bias and brains: Risk aversion and cognitive ability across real and hypothetical settings. *Journal of Risk and Uncertainty*, *46*, 299–320.

Weber, E. U. (1998). Who's afraid of a little risk? New evidence for general risk aversion. In *Decision research from Bayesian approaches to normative systems: Reflections on the contributions of Ward Edwards.* Norwell, MA: Kluwer Academic Publisher.