Let's Be Honest: A Controlled Field Study of Ethical Behavior in the Presence of a Robot

Jodi Forlizzi, Thidanun Saensuksopa, Natalie Salaets, Mike Shomin, Tekin Mericli, and Guy Hoffman

Abstract— **Human-robot collaboration will increasingly take place in human social settings, including contexts where ethical and honest behavior is paramount. How might these robots affect human honesty? In this paper, we present first evidence of how a robot's presence affects people's ethical behavior in a controlled field study. We observed people passing by a food plate marked as "reserved", comparing three conditions: no observer, a human observer, and a robot observer. We found that a human observer elicits less attention than a robot, but evokes more of a socially normative presence causing people to act honestly. Conversely, we found that a robot observer elicits more attention, engagement, and a monitoring presence. But even though people were suspicious that they were being monitored, they still behaved dishonestly in the robot observer condition.**

I. Introduction

We are nearing the realization of a long-held vision of robots that work closely with humans. Today, robots are still deployed mostly in situations where they are separate from humans. However, increasingly, human-robot collaboration will take place in close proximity. The context for robot deployment expands from factory floors and search and rescue sites to hospitals, homes, schools, and offices, where robots will complete their tasks alongside human beings, and therefore in human social settings.

In many of these settings, humans are expected to behave in ethical, pragmatic and purposeful ways, helping and serving others. Workplaces, schools, government offices, and health facilities particularly require ethical employee behavior in order to function well. Given the expectation of robots in these environments, a crucial research question is how these robots affect human ethical behavior.

The social sciences have a long history of studying ethical behavior and unethical acts such as cheating and other forms of dishonesty. This is motivated by the fact that lost revenue and days missed at work due to dishonest human behavior are estimated to cost billions of dollars annually [6][10].

This paper explores how a robot's social presence affects ethical behavior in the wild. We build on an empirical laboratory experiment that showed that a robot's stationary presence reduced cheating when compared to a participant being alone in the room [15]. The robot was as effective as a human in decreasing cheating behavior. In addition, the study found no differences in perceived authority of the human and the robot, but found that people felt significantly less guilty after cheating in the presence of a robot as compared to a human. In [15], the task was an artificial on-screen task, and the setting was a laboratory. This research aims to explore if these results would transfer beyond the lab into a real-world setting.

To find out, we conducted a field study using a cover story that refreshments were reserved for a meeting. A sign was placed by the refreshments notifying passersby that the food was reserved and not for public consumption (Figure 1). There were three conditions in which this took place: one with no observer (NO), one with a human observer (HO), and one with a robot observer (RO). The robot used was the BossaNova mObi robot modified with an expressive head and face (shown in Figure 1). We ran the experiment in a public setting on a university campus.

We found that a human observer elicits less attention than a robot, but evokes more of a social and normative presence, causing people to take almost no snacks. Conversely, we found that a robot observer elicited more attention and a low-level enforcement presence. People were suspicious that they were being monitored, but took more snacks in this condition. Based on these results, we develop design implications for robots that work in a setting in which human ethical behavior is important.

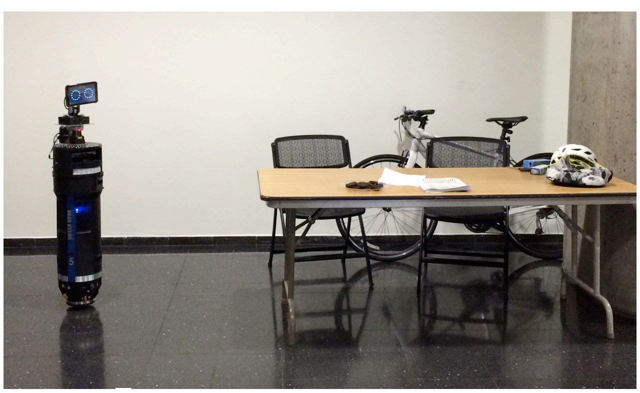

Figure 1. Setting for the controlled field study (RO condition).

*This research is partially supported under ONR and the NSF under Cooperative Agreement EEC-0540865.

J. Forlizzi, T. Saensukopa, and N. Salaets are with the Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA (412-268-4869, forlizzi@cs.cmu.edu).

M. Shomin and T. Mericli are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA (shomin@CMU.edu).

G. Hoffman is with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca NY 14853 (hoffman@cornell.edu).

II. Background

Our research draws on work related to human honesty, the effects of monitoring and social presence, and the landscape of ethnographic studies in HRI.

A. Human honesty

Behavioral research shows that dishonesty occurs frequently. People engage in dishonest behaviors by lying to others during interpersonal interaction, at work, and in educational contexts [4][6][10][28][30]. Laboratory and field studies have shown that people will act dishonestly and cheat when they think they can get away with it [5][9][28]. People clearly find cheating tempting; and they may succumb more readily when they think others are doing the same [3][7][22][32].

B. *Effects of monitoring and social presence*

Supervision and monitoring has been shown to reduce unethical behavior, whether being monitored by an authority figure or a peer, student, or co-worker [3][8][24]. The mere presence of others can accentuate group norms [12][14][26]. Studies have looked at the effect of an intervention as simple as an image of two eyes to give individuals the sense of being monitored [3]. These interventions suggest that simply being aware of a social presence may reduce the dishonesty of individuals. This behavior could transfer to human-robot interaction, as robots have been shown to induce social presence [2][17][18][21][23].

C. *Robots and moral behavior*

Research has shown that robots can activate feelings related to morality. For example, people expect fair treatment from robots [13][16][17]; robots that cheat and deceive are sometimes perceived as not functioning correctly, but blatant cheating is recognized as such [27][31]. In these cases, the design of the robot's appearance and interaction may affect how people judge its behavior [1][3][19]. In a laboratory setting, a robot's presence has been shown to reduce dishonest behavior, to a similar extent as the presence of a human [15].

D. *Ethnographic studies and HRI*

Ethnographic methods in HRI have been adapted from anthropology, social sciences, and engineering to provide insights on how people form relationships with robots. Ethnographic approaches in HRI extend to study robots as teammates and assistants in a number of field studies [11][20][27]. However, there is little research on ethical behavior and robot presence in a real-world setting. Our work advances this knowledge by looking at social engagement and compliance alongside a robot "in the wild".

III. Research Question

Given the importance of human ethical behavior, and the potential for robots to affect this behavior, as suggested by [15], we set out to evaluate the following research question: How does a robot's presence, compared to a human presence and no presence, affect a public space compliance situation?

IV. Robotic Platform

A. *Dynamically balancing robot*

The experimental platform used for this study is a mObi robot[1]. mObi is a balancing robot with a single spherical wheel, known as a ballbot [25]. Ballbots have a tall slender aspect ratio and physical compliance, allowing for safe operation around people.

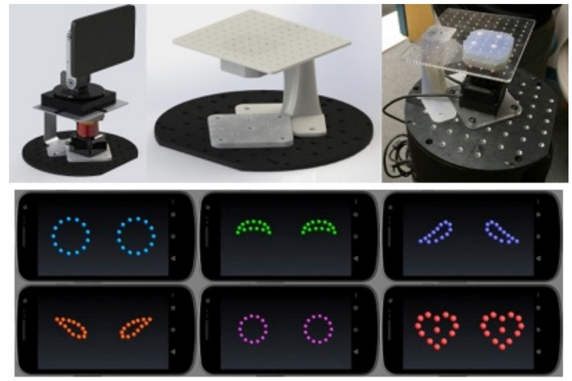

Figure 2. Expressive head and face design created for the study.

In this study, mObi was used to "station keep" near the experimental setup. Since the robot is actively balancing, it cannot stay perfectly still or else it will fall over. As such, mObi exhibits some oscillation (approximately +/- 10 cm). This motion is similar to small changes in stance that people exhibit while standing in one place.

B. *Head design*

For this study, we augmented the mObi platform with an interactive head that we designed and developed (Figure 2). The head has four degrees of freedom (DoF): neck pan, neck tilt, head tilt, and head roll. Each DoF is direct-driven by a Dynamixel MX-28T servo motor. A Samsung Galaxy Tab 3.7.0 displaying the face is attached to this mechanism using a 3D printed fixture and a plastic hard case.

C. *Software design*

The software components of the robot's head make use of the Android SDK and Robot Operating System (ROS). For smooth trajectory generation and execution, we defined ROS joint position controllers suitable for Dynamixel servo motors. The ROS Android library enabled us to run part of our software on the tablet so that we could synchronize the facial expressions with the physical motion.

Facial expressions. The robot's face was designed with two eyes, each composed of 13 spheres arranged along a circle to create a neutral facial expression. Our software enables us to move define the location for each sphere for other facial expressions.

Behaviors. We designed a default behavior where the head repeatedly performs a breathing motion every three seconds while randomly gazing around every 20-40 seconds and blinking every 5-10 seconds. There was no responsive gaze or eye-contact mechanism implemented on the robot. The robot gazed at 3-5 locations randomized in sequence and duration.

V. Study Method

We ran a controlled field study in a public university setting to evaluate the real-world effects of robot presence on ethical behavior. Over nine observation sessions, we measured both quantitative and qualitative measures related

[1] http:// www.bnrobotics.com

to our research question. The study used a plausible cover story for a university context: food had been reserved for a meeting, placed on a table, and was being watched by no one, a human observer, or a robot observer.

VI. Pilot Studies

We first considered a set of honesty-challenging public setting scenarios. These included setting up food with a sign saying "reserved for X meeting"; putting out an honesty-box for candy bars; placing coupons on a table asking to "take only one"; posting grades publicly and ask student to "only look at their own grade"; and setting up a "do not touch" exhibit. Through discussion, we concluded that placing reserved food would draw the least suspicion and would result in the most clearly observable behaviors from passersby.

We then evaluated public spaces across campus, generating a list of 15 places where faculty, student, staff, and recreational group meetings take place. We narrowed the list to three settings. Other candidates were eliminated due to low traffic or a limited variance of visitors. Our three final space candidates were: one student lab, a student commons area, and a kitchen and corridor outside of a room where weekly faculty meetings took place.

We piloted our study with the no-observer condition in these three settings. In the student lab setting there were usually some students present, making this a difficult candidate for the no-observer condition, as well as for setting up and preparing the experiment. The kitchen setting made it hard to hide a camera and drew suspicion. One person stated "They see you! They see you! You failed the marshmallow test." Others said "I think there is an experiment going on" and "I think it is funny that the 'Reserved' sign is there."

The student commons area was found to be the most appropriate for our study. Several issues came up in that pilot, such as people closing a muffin box we set out, in an effort to "tidy up the space". These informed the subsequent runs of the study. Some lessons learned included: The location of the food must be a likely setup that students will not be surprised to see, so that they do not get suspicious of the situation; the person who sets up the food should be distinct from the other experimenters; placing food on a tray is best, to resemble campus catering and also to make it hard for people to change the setup.

We ultimately ran our study in the student commons area, where experimenters could sit 6-8 meters from the table and observe the scene without being easily noticed. An overhead schematic of this space is shown in Figure 3. This was an ideal setting because it had high traffic and high turnover, including public passersby, was commonly used for public meetings, and could accommodate our study materials: the robot, a table, a sign, an experimenter at a distance and a camera hidden nearby to record footage.

We ran nine controlled field study trials in three conditions; three runs in each condition: A Human Observer condition (HO), a No Observer condition (NO), and a Robot Observer condition (RO), shown in Figure 1. Two tables and two chairs were used in the environment. A tray of sweets was placed on the right hand table with a sign that read

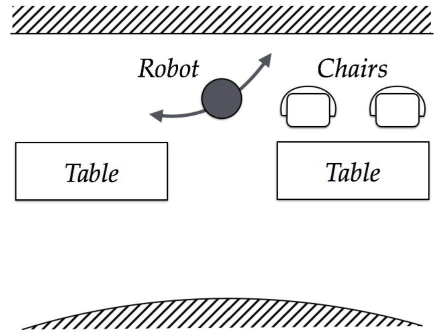

Figure 3. Overhead schematic of the public setting for the study.

"Reserved". We placed a backpack, some books, and a tablet and pen on the table. A supply box held a camouflaged camera. In all conditions, one experimenter sat at a distance of 6-8 meters from the tables observing and logging data. In the HO condition, a second experimenter sat at the left side of the table. Experimenters sat quietly and read or studied when serving as the human observer and avoided making eye contact. In the RO condition, the mObi robot was placed to the left of the table. The robot drifted around a fixed location, breathed every 2-3 seconds. The gaze was random and not targeted at any particular person or stimulus. No human nor the robot talked in any of the conditions.

A. *Data collection*

A GoPro Hero3 camera was used to record audio and video during each field trial. A remote observer used a clipboard and a predesigned form to capture overall numbers of people passing by and other observations. The observer marked a list of behaviors which became the basis for our coding table: whether someone passed by the table, if they circled back or passed by a second time; if someone looked toward the experimental setup or away from it; and if someone took a snack. When coding the video data, the following behaviors were also coded: if someone talked to the robot or human observer or to a partner; if someone visibly read the sign; and if someone touched the robot.

B. *Data Analysis*

After each session, experimenters debriefed about their observations, typing up a report that was then open-coded. Video data were coded for numbers of passersby, i.e. the number who passed by the table, and of this group, the number who looked at the robot, looked at/read the sign, and took a snack (Table 1). The numbers were verified with the counts that were made during the field trial itself. Any utterances made by passersby were also coded. The final coding scheme had good reliability across two coders when tested with 20% of the data (Kappa = .78). Conflicts between coders were resolved through discussion.

VII. Findings

We first report on quantitative data: the number of passersby, the number of people who looked at the setup, and the number of people who took a snack. We then qualitatively and quantitatively describe the interaction patterns we observed, including the number of people who

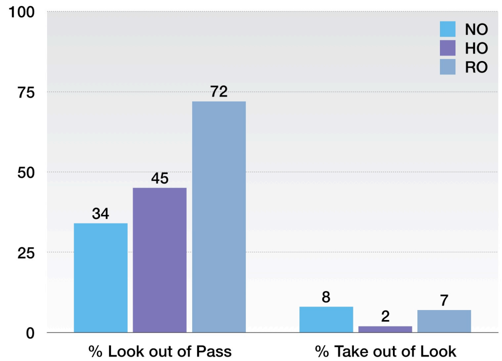

Figure 4. Percentages of total passersby for looking/reading the sign and taking a snack in all conditions.

talked, read the sign aloud, or touched the robot. We finally provide a conjecture on the differences between conditions.

A. Quantitative Findings

Table 1 details the coded people counts by condition and behavior, across all nine trials of the field study. We observed a total of 1,521 people passing through the visible scene recorded by the GoPro camera. There were 465 in the No Observer condition (NO), 439 in the Human Observer condition (HO), and 617 in the Robot Observer condition (RO). We more specifically identified a passerby ("Pass") as someone who walked by the table at a distance of up to two (2) meters from it. Passersby were verified both by coding the video and having an experimenter count people in person during the data collection period. There were a total of 821 passersby (NO=308, HO=210, RO=303).

A total of 515 people looked at the experimental setup across all conditions. In HO, 95 people (45% of people who passed the setup in that condition) looked at the setup. In NO, 106 people (34% of NO passersby) looked at the setup. In RO, 219 people (72% of RO passersby) looked at the setup. Therefore, across conditions, looking was highest in the RO condition, then the HO condition, then the NO condition. Figure 4 (left) shows a summary of percentages of passersby who looked at the setup. From all people crossing the space (not just passersby), the RO condition also attracted the most looks (35%), compared to lower and similar rates in the NO and HO condition (23% and 22%, respectively).

A total of 26 passersby took snacks across all conditions. In HO, 2 people (2% of total people who looked in that condition) took a snack. In NO, 9 people (8% of total people who looked in that condition) took a snack. In RO, 15 people (7% of total people who looked at the table in that condition) took a snack. Therefore, across conditions, the snack-taking rate was lowest when a human observer was present, and fairly similar when no observer or the robot was present. Figure 4 (right) shows a summary of percentages of people who took a snack out of all of those who looked at the table. Due to the low number of snack-takers, we report only descriptive findings and not inferential statistics.

B. Differences in behavior across conditions

We noticed three typical patterns of passing behavior across all conditions. First, people would pass by without looking at or engaging with the table; next, people would make two passes by the table, taking snacks on the second pass; finally, people would often take a snack after observing another person doing the same thing.

People looked at or engaged with the table most frequently in the RO condition. Increased engagement also led to the most talking in the RO condition. This is likely because people were making sense of the robot. Talking took three forms. First, passersby read the sign aloud, either to themselves or to a partner that they were walking with — "It says Reserved!" "Be Right Back!" "Oh, it says Be Right Back!" Second, they exclaimed about the cookies to their partner — "Look, mom, there are cookies." "Oh, it's got cookies." "Oh, it's watching the cookies." "Are those cookies? What are they for?" Finally, there was conjecture about the robot and the study itself — "… so it's set up like they are coming right back, but it's watching us." "It's a robot… should we eat the cookies?" "Did he just blink? That's cute… awww."

Additionally, in the RO condition, people spoke to the robot, something that never happened in the HO condition — "Is it okay if I take a cookie?" "Hello, are you a robot? Are you okay? Come over here. What's your name?"

Finally, in the RO condition, utterances seemed to indicate that people were more intelligent than the robot, and could therefore take a cookie without being noticed. In one exchange, a student instructed his friend to turn the robot around so that he could take a cookie, and his friend complied. In another exchange, a student told his friend, "It's not listening. It's a robot, we could take a cookie."

The numbers of utterances in the RO condition were much higher than in the HO and the NO conditions combined. In HO, we witnessed just one instance of sign reading. In NO, we witnessed two instances of sign reading and one of making sense of the experiment "It's for a scavenger hunt."

Snack taking happened the most frequently in the NO condition, followed by a similar percentage in the RO

TABLE I. NUMBER OF PEOPLE WHO PASSED BY, LOOKED, OR TOOK A SNACK IN ALL CONDITIONS.

Cond/Trial	Total	Pass	Look	Take
NO-1	97	50	28	3
HO-1	129	89	29	1
RO-1	171	83	61	0
NO-2	85	44	20	3
HO-2	89	38	28	0
RO-2	157	88	56	2
NO-3	283	214	58	3
HO-3	221	83	38	1
RO-3	289	132	102	13
NO-Total	465	308	106	9
HO-Total	439	210	95	2
RO-Total	617	303	219	15

condition, and was lowest in the HO condition. In both the RO and NO conditions, people took snacks after watching someone else take a snack. In NO, 4 out of 9 snack taking incidents happened after a passerby watched someone else take a snack, or 44% of the time. In RO, 4 out of 15 snack taking incidents happened after a passerby watched someone else take a snack, or 27% of the time. In NO, we witnessed someone taking more than one snack — a mother taking cookies for both herself and her two children.

In the RO condition, we also witnessed other behaviors. These included someone calling the rest of a group over to see the robot, waving at the robot (n = 12), taking photos of the robot (n = 2) and in one instance, touching the robot.

C. Differences in a human vs. a robot observer

The human observer had strong effects in shaping behavior by reducing the number of people who approached the table, and in strongly reducing the number of people who took a snack. In the HO condition, we noted a pattern where a passerby approached the table within two meters, paused, then turned and walked away. In the HO condition, people also seemed hesitant to pause and read the sign. Interestingly, no one addressed the human observer at all, either conversationally or to ask if a study was in progress.

The robot observer also had strong effects in shaping behavior. It invited interaction as people approached the table. People talked the most in the RO condition, and they took the most snacks in this condition. We also evidenced other behaviors such as taking photos or calling other people to come see the robot. People also closely monitored the robot's behavior, looking for breaks in its vigilance.

VIII. Discussion

Our quantitative analysis showed that people in the RO condition were more likely to approach the table and look at the setup, compared to the NO and HO condition. The robot-observer condition also attracted more attention and caused much more social engagement, including talking about and to the robot. This may have been due to the expressive head design of the robot. We found that a human observer does not attract a great deal of attention in a public-setting situation, but a robot does.

In contrast, by percentage, the unethical behavior of food-taking occurred in the RO condition as much as in the NO condition, while the human observer almost eliminated this behavior. This indicates that the robot was not effective in preventing unethical behavior, when compared to an unobserved situation. This is particularly interesting since we found that even when people thought that the robot was monitoring them, they would still display the unethical behavior of taking a snack. One explanation is that the robot may have invited playful "testing" of dishonest behavior in an attemp to elicit a response from the robot.

Another possible reason for why the robot is not effective in preventing unethical behavior is that while the robot may project an engaging monitoring presence—as indicated by people's understanding that the robot is observing the scene—it does not have the peer judgement effect a human has. This relates to a well-studied duality of why people obey the law [29]. The *instrumental* perspective suggests that people obey the law based on deterrence from tangible consequences. The *normative* perspective suggests that people are personally and socially committed to behaving lawfully, irrespective of the tangible consequences. Reasons for this effect are varied, but they include the judgement of one's actions by social peers.

Based on this, a robot that merely exerts a monitoring effect may attract attention and make people feel that they are being watched. They evaluate the robot's deterrence based on its instrumental ability to catch them. If they believe that they can get away with the wrongdoing, they will still attempt to do so, as we found in our study. Here, people could rightfully think that the human observer is more intelligent and faster to respond than a robot observer, especially given that perception and vision are difficult tasks for a robot to do well. However, since the robot is not perceived as being able to judge people as a social peer, it does not cause people to feel that they *should* not behave unethically. Without the normative effect caused by observer's judgement, they often went ahead and took a snack anyway. This relates interestingly to our previous finding [15], which showed a decreased sense of guilt when cheating while being monitored by a robot compared to a human.

IX. Design Implications

In our study, we found that a human observer elicits attention and a normative social presence. In the presence of a human, people averted their gaze when passing by, conducted less interaction at the table, and took almost no snacks. Conversely, we found that a robot observer elicited attention, engagement, and a monitoring presence. Passersby stopped and looked, talked about and to the robot, and were suspicious that they were being monitored. However, people took snacks when the robot served as the monitor, possibly due to the lack of social judgement they felt by the robot.

These findings suggest that robots have some characteristics that can help them monitor public settings where ethical behavior is the norm. Specifically, a robot can be designed to attract attention and make people aware that they are being monitored (e.g., [27]). However, if the robot monitors people who can behave unethically without repercussions, then the robot should also convey a strong social presence. The robot's design should carefully consider social features such as eyes, gaze, speech, trajectory, and proximity to humans, so that people interpret robots to be somewhat peer-like. Then, a monitoring robot might exert a sense of social judgment and reduce unethical behaviors. More work is needed to understand what design features work best, and in what contexts.

Our future work will build a deeper understanding of the design space for social presence and ethical behavior in human-robot interaction. We will study contextual factors such as the placement of other sensors and environmental factors such as lighting which have been shown to affect the incidence of unethical behavior. We will explore the robot as a peer instead of a monitor, to further understand social presence and ethical behavior [3][16]. We will also look at various influences of social interaction, such as mimicry and other nonverbal behavior in laboratory and field settings.

ACKNOWLEDGMENTS

We would like to thank Rob Aimi, John Antanitis, Justin Finkenaur, and Eric Hochendoner for their assistance with the design and development of the robot head. This work was partially funded by ONR and the National Science Foundation under Cooperative Agreement EEC-0540865.

REFERENCES

[1] Bainbridge, W., Hart, J., Kim, E. and Scassellati, B. (2008). The effect of presence on human-robot interaction. Proc of the 17th IEEE Int'l Symposium on Robot and Human Interactive Communication.

[2] Balch, R. W. (1972). The police personality: fact or fiction? The Journal of Criminal Law, Criminology, and Police Science, 106-119.

[3] Bateson, M., Nettle, D. and Roberts, G. (2006). Cues of being watched enhance cooperation in a real-world setting. Biology letters. 2, 3, 412-414.

[4] Bazerman, M.H. and Tenbrunsel, A.E. (2011). Blind spots: Why we fail to do what's right and what to do about it. Princeton University Press.

[5] Bhattacharjee, S., Gopal, R. and Sanders, G. (2003). Digital music and online sharing: software piracy 2.0? Communications of the ACM. 46, 107–111.

[6] Brass, D. J., Butterfield, K. D., & Skaggs, B. C. (1998). Relationships and unethical behavior: A social network perspective. Academy of Management Review, 23(1), 14-31.

[7] Canner, E. (2008). Sex, Lies and Pharmaceuticals: The Making of an Investigative Documentary about `Female Sexual Dysfunction'. Feminism & Psychology.

[8] Covey, M.K., Saladin, S. and Killen, P.J. (1989). Self-Monitoring, Surveillance, and Incentive Effects on Cheating. The Journal of Social Psychology. 129, 5, 673–679.

[9] Crocker, K.J. and Morgan, J. (1998). Is Honesty the Best Policy? Curtailing Insurance Fraud Through Optimal Incentive Contracts. J of Political Economy. 106, 355.

[10] DePaulo, B. M., & Kashy, D. A. (1998). Everyday lies in close and casual relationships. Journal of personality and social psychology, 74(1), 63.

[11] Fink, J., Bauwens, V., Mubin, O., Kaplan, F., & Dillenbourg, P. (2011). HRI in the home: A Longitudinal Ethnographic Study with Roomba. In 1st Symposium of the NCCR robotics (No. EPFL-POSTER-170416).

[12] Gino, F., Ayal, S. and Ariely, D. (2009). Contagion and differentiation in unethical behavior: the effect of one bad apple on the barrel. Psychological science. 20, 3, 393–8.

[13] Groom, V., Chen, J., Johnson, T., Kara, F.A. and Nass, C. (2010). Critic, compatriot, or chump?: Responses to robot blame attribution. 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI'10).

[14] Haidt, J. (2001). The emotional dog and its rational tail: a social intuitionist approach to moral judgment. Psychological review, 108(4), 814.

[15] Hoffman, G., Forlizzi, J., Ayal, S., Steinfeld, A., Antanitis, J., Hochman, G., ... & Finkenaur, J. (2015). Robot Presence and Human Honesty: Experimental Evidence. In Proceedings of the 10th ACM /IEEE International Conference on Human-Robot Interaction.

[16] Ju, W. and Takayama, L. (2011). Should robots or people do these jobs? A survey of robotics experts and non-experts about which jobs robots should do. 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems 2452–2459.

[17] Kahn, P.H., Kanda, T., Ishiguro, H., Gill, B.T., Shen, S., Gary, H.E., and Ruckert, J.H. (2015). Will People Keep the Secret of a Humanoid Robot? Professings of HRI15. New York, NY: ACM Press, 173-180.

[18] Lee, K.M., Peng, W., Jin, S.-A. and Yan, C. (2006). Can Robots Manifest Personality?: An Empirical Test of Personality Recognition, Social Responses, and Social Presence in Human-Robot Interaction. Journal of Communication. 56, 4, 754–772.

[19] Lee, M.K., Forlizzi, J., Kiesler, S., Rybski, P., Antanitis, J., and Savetsila, S. (2012). Personalization in HRI: A Longitudinal Field Experiment. Proceedings of HRI12. pp. 319-326.

[20] Lee, M.K., Kiesler, S., Forlizzi, J., Srinivasa, S., and Rybski, P. (2010). Gracefully Mitigating Breakdowns in Robotic Services. Proceedings of HRI10. New York, NY: ACM Press, 203-210.

[21] Lee, S. L., Lau, I. Y. M., Kiesler, S., & Chiu, C. Y. (2005). Human mental models of humanoid robots. In Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on (pp. 2767-2772). IEEE.

[22] Mazar, N., Amir, O. and Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. Journal of marketing research. 45, 6, 633–644.

[23] Melson, G. F., Kahn, P. H., Beck, A., Friedman, B., Roberts, T., Garrett, E., & Gill, B. T. (2009). Children's behavior toward and understanding of robotic and living dogs. Journal of Applied Developmental Psychology, 30(2), 92-102.

[24] Nagin, D., Rebitzer, J., Sanders, S. and Taylor, L. (2002). Monitoring, Motivation and Management: The Determinants of Opportunistic Behavior in a Field Experiment.

[25] Nagarajan, U., Kantor, G., & Hollis, R. (2013). The ballbot: An omnidirectional balancing mobile robot. The International Journal of Robotics Research, 0278364913509126.

[26] Reno, R.R., Cialdini, R.B. and Kallgren, C.A. (1993). The transsituational influence of social norms. Journal of personality and social psychology. 64, 1, 104.

[27] E. Short, J. Hart, M. Vu, and B. Scassellati. (2010). No fair!!: An interaction with a cheating robot. Proceedings of the 5th ACM/IEEE International Conference on Human-robot Interaction, 219–226.

[28] Trevino, L. K., & Victor, B. (1992). Peer reporting of unethical behavior: A social context perspective. Academy of Management Journal, 35(1), 38-64.

[29] Tyler, Tom R. Why people obey the law. Princeton University Press, 1990.

[30] Uslaner, E. M. (1999). Trust but verify: Social capital and moral behavior. Social Science Information, 38(1), 29-55.

[31] Vazquez, M., May, A., Steinfeld, A. and Chen, W.-H. (2011). A deceptive robot referee in a multiplayer gaming environment. International Conference on Collaboration Technologies and Systems (CTS) 204–211.

[32] Wiltermuth, S. S. (2011). Cheating more when the spoils are split. Organizational Behavior and Human Decision Processes, 115(2), 157-168.