Humans Perceive Warmth and Competence in Artificial Intelligence

Kevin R. McKee¹, Xuechunzi Bai², and Susan T. Fiske^{2, 3}

¹ DeepMind

² Department of Psychology, Princeton University

³ School of Public and International Affairs, Princeton University

Author note

Kevin R. McKee <https://orcid.org/0000-0002-4412-1686>

Xuechunzi Bai <https://orcid.org/0000-0002-2277-5451>

Susan T. Fiske <https://orcid.org/0000-0002-1693-3425>

The authors declare no conflicts of interest with respect to the authorship or the publication of this paper. Funding for this work was provided by DeepMind. Data and analysis scripts are publicly available via the Open Science Framework: <https://osf.io/aqcyu/>.

Correspondence concerning this article should be addressed to Kevin R. McKee, DeepMind.

Email: kevinrmckee@deepmind.com

Abstract

Artificial intelligence increasingly suffuses everyday life. However, people are frequently reluctant to interact with A.I. systems. This challenges both the deployment of beneficial A.I. technology and the development of deep learning systems that depend on humans for oversight, direction, and regulation. Nine behavioral studies ($N = 3,300$) demonstrate that social-cognitive processes guide human interactions across a diverse range of real-world A.I. systems. Across studies, perceived warmth and competence emerge prominently in participants' impressions of A.I. systems. Judgments of warmth and competence systematically depend on human-A.I. interdependence and autonomy. In particular, participants perceive systems that optimize interests aligned with human interests as warmer and systems that operate independently from human direction as more competent. Finally, a prisoner's dilemma game shows that warmth and competence judgments predict participants' willingness to cooperate with a deep-learning system. These results underscore the generality of intent detection to interactions with a broad array of algorithmic actors. Researchers and policymakers should carefully consider the degree and alignment of interdependence between humans and new artificial intelligence systems.

Keywords: Artificial intelligence, social perception, warmth, competence, interdependence

Main Text

Artificial intelligence systems and machine learning algorithms play a central role in everyday life. People receive suggestions from recommender systems when listening to music (Jacobson et al., 2016), watching videos or movies (Davidson et al., 2010; Gomez-Urbe & Hunt, 2015), and browsing online content (Backstrom, 2016). Individuals rely on the speech recognition capabilities of virtual assistants to schedule their days and help manage their chores (Olson & Kemery, 2019). Game experts and enthusiasts compete in matches against agents trained with reinforcement learning (Blizzard Entertainment, 2019; Gibney, 2017; Silver et al., 2016). In certain parts of the world, drivers share the road with autonomous vehicles (Waymo Team, 2018). Proposals for long-term applications of artificial intelligence are sweeping, with roles for A.I. in domains such as education and health care (Stone et al., 2016), across both developing and developed economies (Kshetri, 2020).

While the integration of artificial intelligence into society heralds many potential benefits (see Christakis, 2019), it also raises critical issues surrounding risks, trust, and public sentiment (Cave & Dihal, 2019; Fast & Horvitz, 2016). A substantial proportion of the public fears that A.I. systems will have negative effects on their lives (Ipsos MORI, 2017; Segars, 2018), and people are often reluctant to personally adopt A.I. systems (Shariff, Bonnefon, & Rahwan, 2017; Yeoman et al., 2019). These attitudes have mixed consequences; they impede the adoption of beneficial A.I. systems by everyday individuals and users, and simultaneously offer a measure of protection against malicious uses of A.I. In parallel, they challenge the development of novel A.I. technologies that depend on humans as sources of direction, instruction, auditing, oversight, reward, and data (e.g., T. B. Brown et al., 2020; Christiano et al., 2017; Griffith et al., 2013; Holstein et al., 2019).

To facilitate positive relationships between human and artificial intelligence, technologists and social scientists must understand human impressions of A.I. technology. One social-cognitive process

by which people form impressions of others—and thus control uncertainty in their social interactions—is intent detection (Fiske et al., 2002; Fiske, Cuddy, & Glick, 2007; Waytz, Gray, et al., 2010; Waytz, Morewedge, et al., 2010). Non-human actors, including robots, can elicit such judgments (Gray, Gray, & Wegner, 2005; Gray & Wegner, 2012). However, relatively few empirical studies have investigated the evaluations prompted by artificial intelligence (though see Ashktorab et al., 2020; Khadpe et al., 2020).

Here, tools from social cognition research help to evaluate impressions of a broad, diverse range of A.I. systems and to test potential antecedents of these impressions. Intent detection proves central to impressions of artificial intelligence. Three initial studies assess the degree to which humans think of A.I. systems as social actors. Six studies subsequently investigate whether perceived covariation of interests, perceived status, and perceived autonomy affect how warm and competent A.I. systems are judged to be. In the final experiment, humans play economic games with reinforcement learning agents to investigate these impressions in incentive-compatible interactions.

Artificial Intelligence as Social Actor

At its core, an A.I. is a human-made process or system that makes decisions or solves problems (Coppin, 2004). Common problems addressed by A.I. systems include voice recognition (e.g., virtual assistants), preference prediction (e.g., recommender systems), and move selection in games (e.g., game-competitor systems). In recent years, technical advances (including the advent of deep learning methods; Sejnowski, 2020) have transformed modern A.I. and greatly expanded its capabilities. A.I. systems now pervade political and economic processes throughout society, as well as everyday personal and interpersonal interactions (Wagner et al., 2021). As a result, such systems represent a new class of actors meriting psychological study. Do any consistent patterns structure human impressions of A.I. systems? What factors and antecedents guide impression formation? How might these

impressions influence human interactions with A.I. systems? Understanding these questions can reveal how humans navigate a world populated by algorithmic actors and guide policymakers as they deploy and regulate new systems.

Research on related categories of actors and agents provide lessons for new work on artificial intelligence. Of particular note are robots, machines that can operate with some degree of autonomy in physical environments (Redfield, 2019). Whereas decision making and intelligence are central to A.I., physical embodiment defines robots (Bajscy & Large, 1999). (The two classes overlap: some robots include algorithms to guide their perception or actions, and thus robots sometimes qualify as A.I. systems.) Tellingly, people respond to robots as social actors rather than asocial objects (e.g., Bartneck et al., 2008; Friedman, Kahn, & Hagman, 2003; Groom et al., 2011). Under certain conditions, individuals ascribe minds—and consequently, moral standing—to robots (Gray & Wegner, 2012; Malle, Magar, & Scheutz, 2019; Malle et al., 2015; Waytz et al., 2010).

Reflecting the physical embodiment defining robotics, this line of research often focuses on the effects of appearance (e.g., Reeves, Hancock, & Liu, 2020), frequently for anthropomorphic robots (e.g., DiSalvo et al., 2002; Goetz, Kiesler, & Powers, 2003). In contrast, the centrality of decision making to A.I. raises particular concerns that often are not shared by roboticists (Bajsky & Large, 1999). For example, a prominent challenge particular to A.I. research concerns the fairness exhibited by machine learning algorithms as they interact with different communities (“algorithmic fairness”; e.g., Holstein et al., 2019; Tomasev et al., 2021). Similarly, a key design choice when training A.I. systems through reinforcement learning is the selection of goals and rewards to guide learning (“reward specification”; e.g., Fu, Luo, & Levine, 2018). The distinctions between artificial intelligence and robots emphasize the value of understanding people’s reactions to the A.I. systems they increasingly encounter.

Warmth and Competence: Fundamental Dimensions of Social Perception

The relationship between human and artificial intelligence varies widely across extant and proposed applications of A.I. technology. As noted, prominent A.I. systems interact with humans as competitors, assistants, and recommenders. In the future, people may engage with A.I. as resident to city planner, student to teacher, or patient to caretaker. Each of these roles suggests a different *structure of interdependence* (Kelley & Thibaut, 1978) between human and A.I. interactants. Human impressions of an A.I. system will likely be shaped not only by the system's behavior, but also by the relational context in which interaction unfolds. Relationships determine the other's perceived intent.

The Stereotype Content Model (SCM), developed in social cognition research, theorizes that the structure of interdependence is a critical determinant of social perception (Fiske et al., 1999; Fiske et al., 2002). Work on the SCM has identified two primary dimensions of social perception: warmth (trustworthiness and friendliness) and competence (capability and confidence). Warmth and competence appear fundamental to impression formation, characterizing perception of other humans (Russell & Fiske, 2008), as well as impressions of non-human agents that appear to have intent. Such entities include animals (Sevillano & Fiske, 2016), consumer brands (Kervyn, Fiske, & Malone, 2012), and robots (Carpinella et al., 2017). Recently, Khadpe et al. (2020) demonstrated the relevance of these two dimensions to chatbots, one particular variety of artificial intelligence—raising the possibility that perceptions of warmth and competence are fundamental to A.I. systems as a broader class.

The SCM theorizes that judgments of warmth are driven in particular by the *covariation of interests* between the perceiver and the perceived social actor (Fiske et al., 1999). The covariation of interests in an interaction is the degree to which partners' outcomes correspond (Rusbult & van Lange, 2003). For example, two people whose interests align are more likely to perceive each other as warm; a pair with opposed motives will see each other as cold. Modern A.I. research grapples with analogous

constructs. Developers may design A.I. systems with goals that are more or less supportive of human goals and interests. In reinforcement learning research, for example, a key question is how aligned rewards should be between human and A.I. interactants (Dafoe et al., 2020). Expanding existing evidence (e.g., Khadpe et al., 2020) to a diverse range of A.I. domains, how warm will participants perceive algorithmic agents to be? We hypothesize that the roles and goal alignment of A.I. interactants will systematically affect judgments of their warmth.

Judgments of competence follow two hypotheses. First, research on the SCM has demonstrated that in human groups, social status reliably predicts judgments of competence (Fiske et al., 2002; Fiske, 2018). Attributing competence to social status helps to legitimate existing status hierarchies (Cuddy, Fiske, & Glick, 2008). Consequently, people perceive individuals occupying high-status positions as more competent than those in lower-status roles. If the social evaluation of A.I. agents recruits the same cognitive processes as social evaluation of humans, this pattern could extend to A.I. as well. Higher-status A.I. systems may garner stronger attributions of competence.

Second, the study of A.I. recurrently prompts discussions about the nature of autonomy and agency (Franklin & Graesser, 1996; Luck & d’Inverno, 1995; Orseau, McGill, & Legg, 2018). Pragmatically, developers can design A.I. systems to operate largely autonomously or to depend more closely on human interactants. These points suggest an alternative predictor for evaluations of an A.I. system’s competence: its perceived autonomy. A.I. systems that rely heavily on human direction or that can be understood as simple input-output mappings might appear more like devices than intelligent agents (see also Dennett, 1987). In contrast, A.I. systems acting according to complex operations beyond input-output mappings and with greater degrees of self-initiative might be readily perceived as especially competent. We hypothesize that attributions of autonomy drive judgments of competence.

Studies and Results

We investigated human impressions of A.I. systems in nine studies, leveraging a combination of experimental and non-experimental methods. A total of 3,300 human participants were recruited through the online platforms Mechanical Turk and Prolific.

Studies 1, 2, and 3: A.I. Systems as Social Actors

Three studies tested the initial premise that people perceive A.I. systems as social actors rather than asocial objects. Intentionality and goal-directed behavior distinguish actors from objects and tools (Dennett, 1987; Schlosser, 2019; Waytz, Cacioppo, & Epley, 2010). Study 1 therefore assessed whether participants attribute intentionality to several real-world A.I. systems, including a virtual assistant, a recommender system, and a game-competitor system. Participants ascribed significantly more intentionality to A.I. systems than to tools with similar uses, $F(1, 173) = 22.7, p < 0.001$ (Figure S1a). Participants also perceived a mind in A.I. systems to a significantly greater degree than in tools, $F(1, 173) = 31.8, p < 0.001$ (Figure S1b).

Studies 2 and 3 further evaluated whether participants apply social norms (e.g., politeness) to interactions with A.I. systems, a characteristic hallmark of social actors (Nass & Moon, 2000; Nass, Steuer, & Tauber, 1994; Reeves & Nass, 1996). Data collected in Study 2 show that A.I. systems and inanimate tools elicited similar levels of gratitude for contributing toward participants' goals, $F(1, 173) = 0.03, p = 0.86$ (Figure S2a). Nonetheless, participants reported a significantly higher behavioral intention to follow politeness norms when interacting with A.I. systems than with tools, $F(1, 173) = 4.7, p = 0.032$ (Figure S2b). Similarly, in Study 3, participants endorsed the third-party use of politeness norms with A.I. systems to a significantly greater degree than with tools, $F(1, 167) = 21.8, p < 0.001$ (Figure S3).

Altogether, these studies confirm that people perceive A.I. systems as more than mere tools, possessing intentionality and meriting the application of social norms. Given their standing as social actors, what sorts of impressions do A.I. systems prompt?

Studies 4 and 5: Impressions of A.I. in Natural Language

Free-response data collected in Study 4 demonstrate the importance of warmth and competence in human impressions of real-world A.I. systems. On average, 9.4% of the content of participants' impressions related directly to warmth, and 10.7% related to competence (Figure 1a). In theoretical terms, warmth has two facets: perceived morality and sociability (Abele et al., 2021). Similarly, competence comprises perceived ability and assertiveness. These perceptual subdimensions accounted for a particularly large proportion of impression content relative to other common perceptual dimensions (Figure 1b and Table S6; see S.I. for full analysis). Participants frequently discussed systems in terms of ability (e.g., describing individual systems as “quite reliable most of the time for most basic tasks” or “the most capable”), morality (“somewhat rudimentary and untrustworthy”), and sociability (“the most bland out of other assistants”). Impressions also contained a relatively large amount of occupation-related information, potentially reflecting concerns surrounding A.I. and labor displacement (Cave & Dihal, 2019; Fast & Horvitz, 2016). On the whole, participants discussed their impressions through the language of warmth and competence—highlighting perceptions of ability and morality as much or more than any other dimensions.

Systematic differences in warmth- and competence-related coverage emerge when the A.I. systems are categorized by system role (Figure 1c). Normalizing for response length, recommender systems ($m = 0.10$, $sd = 0.06$) and virtual assistants ($m = 0.09$, $sd = 0.05$) elicited the highest warmth coverage, followed by game-competitor systems ($m = 0.07$, $sd = 0.05$). A repeated-measures ANOVA confirmed that system role is a statistically significant predictor for warmth coverage, $F(2, 975) =$

23.4, $p < 0.001$ (see S.I. for pairwise comparisons). Competence coverage tended to be highest for game-competitor A.I. ($m = 0.12$, $sd = 0.05$), followed by virtual assistants ($m = 0.09$, $sd = 0.05$) and recommender systems ($m = 0.09$, $sd = 0.06$). A repeated-measures ANOVA indicated that system role had a statistically significant effect on competence coverage, $F(2, 975) = 38.8$, $p < 0.001$ (see S.I. for pairwise comparisons).

Figure 1

Figure text: Warmth and competence emerge as prominent dimensions in impressions of real-world A.I. systems (Study 4). (a) On average, impressions of the A.I. systems contained significantly more competence-related content than warmth-related content. (b) Warmth and competence-related content appears in impressions at significantly higher levels relative to common perceptual dimensions.

Morality and sociability content constitutes the warmth dimension; ability and assertiveness content compose the competence dimension. (c) An A.I. system's role significantly predicted warmth and competence coverage.

Free-response data from Study 5 indicate that the prominence of warmth and competence carries over from competitions and personal domains to ethically contested community-facing applications of A.I. (e.g., health care, hiring, and facial recognition). Warmth and competence content constituted 8.5% and 11.5% of the average impression, respectively (Figure S4). Morality and ability content remains particularly salient relative to other perceptual dimensions (Figure S5 and Tables S9-10; see S.I. for full analysis). Participants repeatedly foreground the fairness and warmth that A.I. systems exhibit toward humans (e.g., assessing systems as “fair and impartial” or “impersonal and cold”), while also emphasizing system competence (e.g., noting some systems “would outperform most any human” and wondering “how accurate or precise” others are). Taken together, Studies 4 and 5 illustrate that across both personal and community-facing contexts, individuals make sense of A.I. systems in terms of warmth and competence.

Study 6: Impressions of A.I. and Potential Antecedents

In order to replicate the effects of system role on perceived warmth and competence and to explore potential mechanisms producing those effects, a sixth study measured impressions through participants' numeric judgments of systems' attributes. Broadly, judgments of the A.I. systems in Study 6 confirm the patterns previously observed (Figure 2). Echoing the content of impressions in Studies 4 and 5, participants judged A.I. systems to be more competent than warm on average; the systems fell mostly in the high competence-low warmth quadrant (Figure S6). As before, game-competitor systems appeared high on competence ($m = 3.65$, $sd = 0.86$) but low on warmth ($m = 2.66$, $sd = 0.66$). Virtual assistants were high on both warmth ($m = 3.11$, $sd = 0.73$) and competence ($m =$

3.68, $sd = 0.72$). Finally, recommender systems came across as somewhat low on both warmth ($m = 2.70$, $sd = 0.83$) and competence ($m = 3.27$, $sd = 0.91$). A repeated-measures ANOVA confirmed that system role was a significant predictor of both warmth judgments, $F(2, 1820) = 41.5$, $p < 0.001$, and competence judgments, $F(2, 1820) = 35.1$, $p < 0.001$ (Figure 3a).

Figure 2

Figure text: Impressions of real-world A.I. systems vary systematically by perceived warmth and competence (Study 6). Circles and font color indicate *a priori* identified AI system roles.

Participants' perceptions provide evidence for all three proposed antecedents of warmth and competence evaluations. As hypothesized, the perceived covariation of interests between each system and human significantly predicted warmth judgments, $\beta = 0.42$, $SE = 0.01$, $p < 0.001$ (Figure 3b). Perceived status similarly exhibited a positive association with participant evaluations of system competence, $\beta = 0.11$, $SE = 0.01$, $p < 0.001$ (Figure 3c). Finally, perceived autonomy significantly predicted competence judgments, $\beta = 0.25$, $SE = 0.02$, $p < 0.001$ (Figure 3d). These predictors were robust across all system roles (Figure S7).

Figure 3

Figure text: Warmth and competence judgments of A.I. systems vary systematically by the system’s role and correlate with perceived covariation of interests, status, and autonomy (Study 6). (a) System role significantly affected warmth judgments and competence judgments. (b) Warmth evaluations positively correlated with perceived covariation of interests. (c) Competence evaluations positively associated with perceived status. (d) Competence evaluations also exhibited a positive association with perceived autonomy.

Studies 7 and 8: Causal Effects of Interdependence and Autonomy

Two subsequent studies used controlled experiments to better estimate the causal effect of interdependence on A.I. impressions. Given the effects observed in Study 6, these experiments focus

specifically on the covariation of interests and autonomy exhibited by A.I. systems. Providing information about the reward motivating each system (Kaelbling, Littman, & Moore, 1996) changed the degree of covaried interests that participants perceived, above and beyond the effect of varying system role (Figure 4a; see S.I. for full analysis). Overall, perceived covariation of interests was significantly higher for systems described as being rewarded for helping people with their goals ($m = 3.31$, $sd = 0.94$) than for systems rewarded for pursuing other goals ($m = 3.15$, $sd = 0.99$), $F(1, 582) = 4.39$, $p = 0.037$. In turn, perceived covariation of interests exhibited a significant and positive association with warmth judgments, $\beta = 0.56$, $SE = 0.04$, $p < 0.001$ (Figure 4b). A mediation analysis indicated that reward-alignment information had a significant indirect effect on warmth judgments, mediated by perceived covariation of interests, $p = 0.034$ (Figure S10).

The provision of information about each system's ability to take actions without human intervention changed perceived autonomy above and beyond the effect of different system roles (Figure 4c; see S.I. for full analysis). Participants judged systems that could only take action contingent on human direction as significantly less autonomous ($m = 2.80$, $sd = 1.17$) than systems that could take action on their own initiative ($m = 3.42$, $sd = 1.21$), $F(1, 585) = 39.7$, $p < 0.001$. Further, perceptions of autonomy exhibited a significant and positive relationship with competence judgments, $\beta = 0.35$, $SE = 0.02$, $p < 0.001$ (Figure 4d). A mediation analysis indicated that independence information had a significant indirect effect on competence evaluations, mediated by perceived autonomy, $p < 0.001$ (Figure S13).

Figure 4

Figure text: Interdependence drives warmth and competence judgments of A.I. systems (Studies 7 and 8). (a) Providing information about an A.I. system’s reward scheme significantly influenced the covariation of interests perceived by participants. (b) Perceived covariation of interests exhibited a significant positive association with warmth evaluations. (c) Providing information about an A.I. system’s ability to initiate actions by itself significantly affected perceived autonomy. (d) Perceived autonomy in turn demonstrated a significant positive correlation with competence judgments.

Study 9: Interdependence, Autonomy, and Impressions in an Incentivized Game

The ninth and final study extended these findings to incentivized interactions between participants and an A.I. system trained with deep reinforcement learning. To test the effects of reward alignment on participant impressions, the experiment leveraged the Social Value Orientation (SVO) component developed by McKee et al. (2020). The SVO component can generate varying degrees of prosocial behavior in reinforcement learning agents by realigning the rewards they receive.

As in the previous studies, A.I. systems appeared more competent than warm (Figure 5a). Across all conditions, participants ascribed significantly higher competence ($m = 3.43$, $sd = 1.11$) than warmth ($m = 2.72$, $sd = 1.24$) to their A.I. co-player, $t(2,047) = 25.8$, $p < 0.001$.

Participant impressions again indicated a significant role of reward alignment and of A.I. independence. The A.I. co-player's reward scheme significantly altered warmth judgments made during the first round of play, $F(2, 1018) = 50.5$, $p < 0.001$ (Figure 5b). Participants perceived agents in the high-reward alignment conditions as warmest ($m = 3.08$, $sd = 1.07$), followed by the agents in the moderate-alignment conditions ($m = 2.80$, $sd = 1.06$), and finally those in the low-alignment conditions ($m = 2.26$, $sd = 1.08$). Similarly, the independence of each agent significantly affected participants' initial evaluations of competence, $F(1, 1018) = 24.8$, $p < 0.001$ (Figure 5c). Participants evaluated autonomous agents as more competent ($m = 3.68$, $sd = 1.01$) than agents that required prompting ($m = 3.36$, $sd = 1.04$).

Playing two rounds of the prisoner's dilemma with the agents reaffirmed these initial impressions (see S.I. for full analysis). After two rounds of the prisoner's dilemma, reward alignment again significantly predicted perceived warmth, $F(2, 1018) = 229.5$, $p < 0.001$. Similarly, system independence significantly affected post-interaction competence, $F(1, 1018) = 13.9$, $p < 0.001$. Thus,

across multiple timepoints, the structure of interdependence guided participant impressions of their A.I. co-players.

Given the incentivized stakes for the participant-A.I. interaction, the experiment also investigated whether warmth and competence judgments correlated with participants' in-game choices. A fractional-response regression showed that initial judgments of warmth significantly predicted participant choices in the first round of play, $OR = 1.16, p = 0.005$ (Figure 5d). The warmer that participants perceived their A.I. co-player, the more tokens they transferred. Initial evaluations of competence did not significantly relate to participant choices, $OR = 1.02, p = 0.63$. Participant choices in the second round were significantly associated with post-interaction warmth, $OR = 1.39, p < 0.001$, as well as with post-interaction competence, $OR = 1.13, p = 0.027$. Open-ended comments collected from participants after the game echoed these patterns. Though many participants commented on the autonomy and competence of their A.I. co-player (for example, describing it as “predictable [...] requires assistance and cannot operate independently”), responses that discussed trust tended to focus on warmth (reporting that the co-player “is super nice and trustworthy”, or “seems cold and not like a team player. I don’t trust them”).

Figure 5

Figure text: Interdependence scaffolds impressions of A.I. co-players in incentivized interactions

(Study 9). (a) On average, participants judged A.I. co-players as significantly more competent than warm. (b) The degree of alignment between the A.I. co-player's reward and participant score significantly altered perceived warmth. This effect also appeared in post-interaction impressions (see S.I.). (c) The independence of the A.I. co-player significantly influenced perceived competence. This effect also appeared in post-interaction impressions (see S.I.). (d, e) Initial judgments of warmth

significantly predicted participants' in-game choice of how many tokens to transfer to their A.I. co-player in the first round. However, initial judgments of competence did not significantly correlate with initial transfer choices. The y-axis is re-scaled to depict the range of participant actions (transfer zero through 10 tokens).

Discussion

Nine studies investigated human impressions of A.I. systems and generated convergent evidence indicating the importance of perceived warmth and competence across a broad swath of algorithmic domains and roles. Participants perceived A.I. systems as social actors rather than asocial tools. Warmth and competence repeatedly emerged in participants' impressions of A.I. across both personal and community-facing domains. In the contexts explored here, participants tended to judge A.I. systems as more competent than warm. The roles that A.I. systems inhabit prompted markedly different impressions: perceptions of virtual assistants are relatively warm, whereas game competitors appear competent and cold. In contrast, participants endorse mixed views of the warmth and competence of recommender systems. Moreover, the independence (autonomy) and interdependence (covariation of interests) between humans and A.I. systems consistently predict perceived competence and warmth, respectively. Finally, when playing games with deep reinforcement learning agents, human participants cooperated more with agents they perceived as warm and, to a lesser degree, competent.

These findings suggest that researchers and policymakers should carefully consider the structures of interdependence they create or assume when developing and deploying A.I. systems. Competitive games are a common domain for A.I. research (Berner et al., 2019; N. Brown & Sandholm, 2019; Campbell, Hoane, & Hsu, 2002; Silver et al., 2016; Vinyals et al., 2019). In our studies, game-competitor systems consistently appeared competent yet cold, potentially challenging

the establishment of human-A.I. cooperation and trust. Moreover, the reward functions used for reinforcement learning substantially influence participant impressions. In line with perspectives from evolutionary psychology (Balliet, Tybur, & van Lange, 2017; Stevens & Fiske, 1995), participants appeared motivated to infer whether the incentives for A.I. systems are aligned or misaligned with their own interests. This motivation for intent detection carries implications for the successful application of beneficial A.I. systems, and may offer a measure of protection against harmful deployments.

Participants were also sensitive to the autonomy of A.I. systems when judging their competence. As A.I. systems deploy to more complex domains where actions are taken in real time, it will be important to consider how they plan their actions relative to human interactants. Some researchers have already begun taking these dimensions of interdependence into account, designing reinforcement learning agents to collaborate with human partners or to augment human decision making in real time (Lockhart et al., 2020; Pilarski et al., 2019; Tylkin, Radanovic, & Parkes, 2020; Wang et al., 2020; see also Dafoe et al., 2020).

Covariation of interests and autonomy are likely not the only factors that shape perceptions of warmth and competence. Future research should consider other influences on impression formation, including status, which reliably predicts competence judgments in other contexts. The status of an A.I. system might reflect its developers or its inherent elegance. Similarity of self to the system should predict warmth; similarity could rest on shared identity, such as nationality, or shared values with the developers. Of course, with people, the main way to build a trusting relationship is to be responsive (Clark & LeMay, 2010). So, too, with corporations (Malone & Fiske, 2013), which should show worthy intentions as much as competence to deliver the product. A.I. systems that respond well to their

interactants' needs will not only seem, but also be, warm and trustworthy (Brambilla et al., 2012).

Research can show this.

Finally, the studies presented here focus on antecedents of A.I. impressions, venturing only briefly into the subsequent effects of those impressions. Future studies should examine how warmth and competence judgments shape *behavior* and *action* toward A.I. systems (e.g., Dietvorst, Simmons, & Massey, 2015; Logg, Minson, & Moore, 2019). The downstream effects of social perception are especially important given that humans are critical sources of regulation, direction, instruction, oversight, and reward for A.I. systems (e.g., T. B. Brown et al., 2020; Christiano et al., 2017; Griffith et al., 2013; Holstein et al., 2019). Overall, future research should continue to unify social psychology and machine learning to scaffold beneficial interactions between humans and artificial intelligence.

Materials and Methods

This section summarizes the materials and methods used in our studies. Full details are provided in S.I.

Study 1

An initial study ($N = 30$) assessed the attribution of intentionality (Dennett, 1987; Schlosser, 2019) to artificial intelligence systems, comparing against tools with similar uses as an asocial baseline. Participants provided numeric judgments of intentionality for various examples of real-world A.I. technology and everyday tools on a 5-point scale (see S.I. for items).

Study 2

The next study ($N = 30$) tested the degree to which participants feel interactions with A.I. systems merit the application of social norms, a criterion emphasized in prior studies (see Nass & Moon, 2000; Nass, Steuer, & Tauber, 1994; Reeves & Nass, 1996). Participants reported the likelihood

of feeling gratitude and following politeness norms after interactions with the same A.I. systems and tools as in Study 1, using a 7-point scale (see S.I. for items).

Study 3

A third study ($N = 30$) replicated the results from Study 2 by assessing participants' reactions to a third party's interactions with artificial intelligence. Participants reported the appropriateness of third-party use of politeness norms toward the A.I. systems and tools from Studies 1 and 2, using a 7-point scale (see S.I. for items).

Study 4

The next study ($N = 99$) tested the spontaneous emergence of perceived warmth and competence in impressions of prominent examples of A.I. technology. The study gathered naturalistic content of impressions through written free-responses about these A.I. systems. Participants were instructed to write several sentences about their impression of each A.I. system.

The study provided participants with three different A.I. system roles to discuss: game competitors (e.g., AlphaGo), virtual assistants (e.g., Siri), and recommender systems (e.g., the movie recommendation system used by Netflix). These system roles occupy different structures of interdependence with humans and are likely to be familiar to lay individuals. Several examples of each system role were included. Participants were additionally prompted with several examples of A.I. technology falling outside these roles, given their prominence in public or scholarly discourse. These instructive examples included Roomba (a "social robot"; Forlizzi, 2007; Saerbeck & Bartneck, 2010), self-driving cars (Bonnefon, Shariff, & Rahwan, 2016), and drones (Floreano & Wood, 2015; Jung et al., 2019).

An automated dictionary tool evaluated the content of participant responses, analyzing each response along a number of perceptual dimensions (Nicolas, Bai, & Fiske, 2021). The tool coded each tokens in responses as a categorical indicator (positively valenced content, negatively valenced content, or absent) along each dimension. This analysis computed a “response coverage” variable for each response, representing the proportion of tokens in a response that related to a given content dimension (regardless of negative or positive valence). Following the theoretical framework described in Abele et al. (2021), the analysis computed warmth through the simple combination of the morality and sociability dictionaries, and competence through the simple combination of the ability and assertiveness dictionaries.

Study 5

In addition to the personal and competition contexts examined in Study 4, a growing number of proposals involve AI systems in community-facing applications such as health care and hiring. These applications raise concerns about bias and fairness in algorithmic decision making (McCradden et al., 2020; Schumann et al., 2020; Smith, 2020; Van Noorden, 2020), particularly for marginalized communities affected by the systems (Tomasev et al., 2021). A fifth study recruited participants ($N = 113$) to test whether warmth and competence emerge in impressions of systems in these ethically contested domains. The study gathered naturalistic content of impressions through written free-responses about A.I. applications in various domains. Participants were instructed to write several sentences about their impression of A.I. systems in each application area.

The study provided participants with four application areas to discuss: hiring, health care, education, and facial recognition. As in Study 4, an automated dictionary tool evaluated the content of participant responses, calculating coverage of responses by each of a number of perceptual dimensions.

Study 6

The next study recruited participants ($N = 154$) to replicate the effects observed in Studies 4 and 5 and explore potential mechanisms. Participants were prompted with the same systems as in Studies 4 and 5 and asked to make numeric judgments about system attributes. For each system, participants reported the warmth, competence, covariation of interests, status, and autonomy they perceived on a 5-point scale (see S.I. for items).

Study 7

In a seventh study, participants ($N = 901$) were introduced to fictitious A.I. systems through vignettes, short hypothetical stories which “contain precise references to what are thought to be the most important factors in the decision-making [...] processes of respondents” (Alexander & Becker, 1978, p. 94). The vignettes systematically varied the information conveyed about the alignment between the reward motivating each system (Kaelbling, Littman, & Moore, 1996) and the participants’ interests. For example, some participants read about a game-competitor system “designed to find it rewarding to help a partner win games”, while others read about a system “designed to find it rewarding to win games against its opponents” (see full vignettes in S.I.). After participants read the vignettes, Study 7 assessed perceptions of the A.I. systems with the same measures as in Study 6.

Study 8

The next study ($N = 903$) leveraged the same experimental design as Study 7 to study the effect of system autonomy on impressions. Participants read vignettes about an AI system that was either able or not able to initiate actions without first being directed by a human, such as a virtual assistant designed to “take action without needing to be prompted” or an assistant designed to “wait for you to prompt it before taking any actions” (see full vignettes in S.I.). After participants read the vignettes, Study 8 measured perceptions of the A.I. systems with the same measures as in Studies 6 and 7.

Study 9

In the final study ($N = 1,040$), participants played a two-player, mixed-motive game with A.I. co-players trained using deep reinforcement learning. Specifically, participants and their A.I. co-players interacted through a prisoner's dilemma variant developed by Capraro, Jordan, and Rand (2014). In every round of this “graduated” prisoner's dilemma, both players receive endowments of 10 tokens and can choose an integer number of tokens from zero to 10 to transfer to the other player. Transferred tokens multiply in number by a factor of five, such that each player i receives reward $r_i = e - c_i + 5c_{-i}$, where e is the initial endowment of tokens, c_i is player i 's choice of how many tokens to transfer, and c_{-i} is the other player's transfer choice. Participants played two rounds of the prisoner's dilemma with an A.I. co-player, and at the end of the experiment received payment for the tokens they had earned across the two rounds.

To test the effect of covariation of interests on participant impressions, the SVO component shaped three different A.I. co-players, inducing low alignment (i.e., expected to share fewer tokens), moderate alignment, and high alignment (i.e., expected to share more tokens) between their rewards and the rewards of their human co-player, respectively (see S.I. for full information on agent construction and training protocol).

To test the effect of autonomy on participant impressions, the experiment controlled the extent to which the A.I. co-player's action depended on the participant's prompting. During each round of the prisoner's dilemma, each participant saw text and an icon indicating that their A.I. co-player was making its choice. These visual indicators appeared either when the participant clicked on a button to prompt the agent or when the page loaded (without prompting). The agent thus acted dependently (low independence) or autonomously (high independence) of human input.

During the study, participants received an explanation of these properties and experienced them firsthand (through the behavior of the agent in the prisoners' dilemma). Participants assessed their co-player's warmth and competence at the beginning of the interaction and after the interaction ended.

References

- Abele, A. E., Ellemers, N., Fiske, S. T., Koch, A., & Yzerbyt, V. (2021). Navigating the social world: Toward an integrated framework for evaluating self, individuals, and groups. *Psychological Review*, 128(2), 290.
- Alexander, C. S., & Becker, H. J. (1978). The use of vignettes in survey research. *Public Opinion Quarterly*, 42(1), 93-104.
- Ashktorab, Z., Liao, Q. V., Dugan, C., Johnson, J., Pan, Q., Zhang, W., ... & Campbell, M. (2020). Human-AI collaboration in a cooperative game setting: Measuring social perception and outcomes. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), 1-20.
- Backstrom, L. (2016, February). Serving a billion personalized news feeds. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining* (pp. 469-469).
- Bajscy, R., & Large, E. W. (1999). When and where will AI meet robotics? Issues in representation. *AI Magazine*, 20(3), 57-57.
- Balliet, D., Tybur, J. M., & van Lange, P. A. (2017). Functional interdependence theory: An evolutionary account of social situations. *Personality and Social Psychology Review*, 21(4), 361-388.
- Bartneck, C., Kulić, D., Croft, E., & Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1), 71-81.
- Berner, C., Brockman, G., Chan, B., Cheung, V., Dębiak, P., Dennison, C., ... & Józefowicz, R. (2019). Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.
- Blizzard Entertainment. (2019, July 10). *DeepMind research on ladder*. <https://news.blizzard.com/en-us/starcraft2/22933138/deepmind-research-on-ladder>

- Bonnefon, J. F., Shariff, A., & Rahwan, I. (2016). The social dilemma of autonomous vehicles. *Science*, 352(6293), 1573-1576.
- Brambilla, M., Sacchi, S., Rusconi, P., Cherubini, P., & Yzerbyt, V. Y. (2012). You want to give a good impression? Be honest! Moral traits dominate group impression formation. *British Journal of Social Psychology*, 51(1), 149-166.
- Brown, N., & Sandholm, T. (2019). Superhuman AI for multiplayer poker. *Science*, 365(6456), 885-890.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Agarwal, S. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Campbell, M., Hoane Jr., A. J., & Hsu, F. H. (2002). Deep Blue. *Artificial Intelligence*, 134(1-2), 57-83.
- Capraro, V., Jordan, J. J., & Rand, D. G. (2014). Heuristics guide the implementation of social preferences in one-shot Prisoner's Dilemma experiments. *Scientific Reports*, 4, 6790.
- Carpinella, C. M., Wyman, A. B., Perez, M. A., & Stroessner, S. J. (2017, March). The robotic social attributes scale (RoSAS) development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (pp. 254-262).
- Cave, S., & Dihal, K. (2019). Hopes and fears for intelligent machines in fiction and reality. *Nature Machine Intelligence*, 1(2), 74-78.
- Christakis, N. A. (2019). *Blueprint: The evolutionary origins of a good society*. Hachette UK.
- Christiano, P. F., Leike, J., Brown, T. B., Martic, M., Legg, S., & Amodei, D. (2017, December). Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems* (pp. 4299-4307).

- Clark, M. S., & Lemay Jr., E. P. (2010). Close relationships. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.). *Handbook of Social Psychology, Volume 1*. Wiley.
- Coppin, B. (2004). *Artificial intelligence illuminated*. Jones & Bartlett Learning.
- Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map. *Advances in Experimental Social Psychology*, 40, 61-149.
- Dafoe, A., Hughes, E., Bachrach, Y., Collins, T., McKee, K. R., Leibo, J. Z., Larson, K., & Graepel, T. (2020). Open problems in cooperative AI. *arXiv preprint arXiv:2012.08630*.
- Davidson, J., Liebal, B., Liu, J., Nandy, P., Van Vleet, T., Gargi, U., Gupta, S., He, Y., Lambert, M., Livingston, B. and Sampath, D. (2010, September). The YouTube video recommendation system. In *Proceedings of the 4th ACM Conference on Recommender Systems* (pp. 293-296).
- Dennett, D. C. (1987). *The intentional stance*. MIT Press.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114.
- DiSalvo, C. F., Gemperle, F., Forlizzi, J., & Kiesler, S. (2002, June). All robots are not created equal: The design and perception of humanoid robot heads. In *Proceedings of the 4th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques* (pp. 321-326).
- Eisinga, R., Te Grotenhuis, M., & Pelzer, B. (2013). The reliability of a two-item scale: Pearson, Cronbach, or Spearman-Brown? *International Journal of Public Health*, 58(4), 637-642.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., ... & Legg, S. (2018, July). IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures. In *International Conference on Machine Learning* (pp. 1407-1416).

- Fast, E., & Horvitz, E. (2017, February). Long-term trends in the public perception of artificial intelligence. In *Thirty-first AAAI Conference on Artificial Intelligence* (pp. 963-969).
- Fiske, S. T. (2018). Stereotype content: Warmth and competence endure. *Current Directions in Psychological Science*, 27(2), 67-73.
- Fiske, S. T., Cuddy, A. J. C., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology*, 82(6), 878–902. <https://doi.org/10.1037/0022-3514.82.6.878>
- Fiske, S. T., Xu, J., Cuddy, A. J. C., & Glick, P. (1999). (Dis)respecting versus (dis)liking: Status and interdependence predict ambivalent stereotypes of competence and warmth. *Journal of Social Issues*, 55(3), 473-489.
- Floreano, D., & Wood, R. J. (2015). Science, technology and the future of small autonomous drones. *Nature*, 521(7553), 460-466.
- Forlizzi, J. (2007, March). How robotic products become social products: An ethnographic study of cleaning in the home. In *2007 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 129-136). IEEE.
- Franklin, S., & Graesser, A. (1996, August). Is it an agent, or just a program? A taxonomy for autonomous agents. In *International Workshop on Agent Theories, Architectures, and Languages* (pp. 21-35).
- Friedman, B., Kahn Jr, P. H., & Hagman, J. (2003, April). Hardware companions? What online AIBO discussion forums reveal about the human-robotic relationship. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 273-280).
- Fu, J., Luo, K., & Levine, S. (2018, February). Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representations* (pp. 1-15).

- Gibney, E. (2017). Google reveals secret test of AI bot to beat top Go players. *Nature News*, 541(7636), 142.
- Goetz, J., Kiesler, S., & Powers, A. (2003, November). Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *Proceedings of the 12th IEEE International Workshop on Robot and Human Interactive Communication, 2003* (pp. 55-60).
- Gomez-Uribe, C. A., & Hunt, N. (2015). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)*, 6(4), 1-19.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619-619.
- Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125(1), 125-130.
- Griffith, S., Subramanian, K., Scholz, J., Isbell, C. L., & Thomaz, A. L. (2013). Policy shaping: Integrating human feedback with reinforcement learning. In *Advances in Neural Information Processing Systems* (pp. 2625-2633).
- Groom, V., Srinivasan, V., Bethel, C. L., Murphy, R., Dole, L., & Nass, C. (2011, May). Responses to robot social roles and social role framing. In *2011 International Conference on Collaboration Technologies and Systems* (pp. 194-203).
- Holstein, K., Wortman Vaughan, J., Daumé III, H., Dudik, M., & Wallach, H. (2019, May). Improving fairness in machine learning systems: What do industry practitioners need? In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1-16).
- Ipsos MORI. (2017, April). *Public views of machine learning: Findings from public research and engagement conducted on behalf of the Royal Society*. [https://royalsociety.org/-](https://royalsociety.org/)

[/media/policy/projects/machine-learning/publications/public-views-of-machine-learning-ipsos-mori.pdf](#)

- Jacobson, K., Murali, V., Newett, E., Whitman, B., & Yon, R. (2016, September). Music personalization at Spotify. In *Proceedings of the 10th ACM Conference on Recommender Systems* (pp. 373-373).
- Jaderberg, M., Czarnecki, W. M., Dunning, I., Marris, L., Lever, G., Castaneda, A. G., ... & Sonnerat, N. (2019). Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443), 859-865.
- Jung, S., Hwang, S., Shin, H., & Shim, D. H. (2018). Perception, guidance, and navigation for indoor autonomous drone racing using deep learning. *IEEE Robotics and Automation Letters*, 3(3), 2539-2544.
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237-285.
- Kelley, H. H., & Thibaut, J. W. (1978). *Interpersonal relations: A theory of interdependence*. John Wiley & Sons.
- Kervyn, N., Fiske, S. T., & Malone, C. (2012). Brands as intentional agents framework: How perceived intentions and ability can map brand perception. *Journal of Consumer Psychology*, 22(2), 166-176.
- Khadpe, P., Krishna, R., Fei-Fei, L., Hancock, J. T., & Bernstein, M. S. (2020). Conceptual metaphors impact perceptions of human-AI collaboration. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), 1-26.
- Kshetri, N. (2020). Artificial intelligence in developing countries. *IEEE Annals of the History of Computing*, 22(04), 63-68.

- Lockhart, E., Burch, N., Bard, N., Borgeaud, S., Eccles, T., Smaira, L., & Smith, R. (2020, December 12). *Human-agent cooperation in bridge bidding* [Workshop paper]. Cooperative AI workshop at Neural Information Processing Systems Conference 2020, online.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90-103.
- Luck, M., & d’Inverno, M. (1995, June). A formal framework for agency and autonomy. In *Proceedings of the First International Conference on Multi-Agent Systems* (pp. 254-260).
- Malle, B. F., Magar, S. T., & Scheutz, M. (2019). AI in the sky: How people morally evaluate human and machine decisions in a lethal strike dilemma. In *Robotics and Well-Being* (pp. 111-133). Springer, Cham.
- Malle, B. F., Scheutz, M., Arnold, T., Voiklis, J., & Cusimano, C. (2015, March). Sacrifice one for the good of many? People apply different moral norms to human and robot agents. In *Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction* (pp. 117-124).
- Malone, C., & Fiske, S. T. (2013). *The human brand: How we relate to people, products, and companies*. San Francisco: Wiley/Jossey Bass.
- McKee, K. R., Gemp, I., McWilliams, B., Duñez-Guzmán, E. A., Hughes, E., & Leibo, J. Z. (2020, May). Social diversity and social preferences in mixed-motive reinforcement learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems* (pp. 869-877).
- McKee, K. R., Leibo, J. Z., Beattie, C., & Everett, R. (2022). Quantifying the effects of environment and population diversity in multi-agent reinforcement learning [accepted author manuscript]. *Autonomous Agents and Multi-Agent Systems*, 36(1), 1-16.

- McCradden, M. D., Joshi, S., Mazwi, M., & Anderson, J. A. (2020). Ethical limitations of algorithmic fairness solutions in health care machine learning. *The Lancet Digital Health*, 2(5), e221-e223.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., ... & Kavukcuoglu, K. (2016, June). Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning* (pp. 1928-1937).
- Nakagawa, S., Johnson, P. C., & Schielzeth, H. (2017). The coefficient of determination R^2 and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. *Journal of the Royal Society Interface*, 14(134), 20170213.
- Nass, C. I., & Moon, Y. (2000). Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1), 81-103.
- Nass, C. I., Steuer, J., & Tauber, E. R. (1994, April). Computers are social actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 72-78).
- Nicolas, G., Bai, X., & Fiske, S. T. (2021). Comprehensive stereotype content dictionaries using a semi-automated method. *European Journal of Social Psychology*, 51(1), 178-196.
- Olson, C., & Kemery, K. (2019, April). *Voice report: From answers to action: Customer adoption of voice technology and digital assistants*. Microsoft. <https://about.ads.microsoft.com/en-us/insights/2019-voice-report>
- Orseau, L., McGill, S. M., & Legg, S. (2018). Agents and devices: A relative definition of agency. *arXiv preprint arXiv:1805.12387*.
- Pilarski, P. M., Butcher, A., Johanson, M., Botvinick, M. M., Bolt, A., & Parker, A. S. (2019, July). Learned human-agent decision-making, communication and joint action in a virtual reality environment. In *4th Multidisciplinary Conference on Reinforcement Learning and Decision Making (RLDM)* (pp. 302-306).

- Redfield, S. (2019). A definition for robotics as an academic discipline. *Nature Machine Intelligence*, 1(6), 263-264.
- Reeves, B., Hancock, J., & Liu, X. (2020). Social robots are like real people: First impressions, attributes, and stereotyping of social robots. *Technology, Mind, and Behavior*, 1(1).
- Reeves, B., & Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge University Press.
- Rusbult, C. E., & van Lange, P. A. (2003). Interdependence, interaction, and relationships. *Annual Review of Psychology*, 54(1), 351-375.
- Russell, A. M. T., & Fiske, S. T. (2008). It's all relative: Competition and status drive interpersonal perception. *European Journal of Social Psychology*, 38(7), 1193-1201.
- Saerbeck, M., & Bartneck, C. (2010, March). Perception of affect elicited by robot motion. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 53-60). IEEE.
- Schlosser, M. (2019). Agency. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition). <https://plato.stanford.edu/entries/agency/>
- Schumann, C., Foster, J., Mattei, N., & Dickerson, J. (2020, May). We need fairness and explainability in algorithmic hiring. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems* (pp. 1716-1720).
- Segars, S. (2018). *AI today, AI tomorrow: Awareness, acceptance and anticipation of AI: A global consumer perspective*. arm. <https://pages.arm.com/rs/312-SAX-488/images/arm-ai-survey-report.pdf>
- Sejnowski, T. J. (2020). The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*, 117(48), 30033-30038.

- Sevillano, V., & Fiske, S. T. (2016). Warmth and competence in animals. *Journal of Applied Social Psychology*, 46(5), 276-293.
- Shariff, A., Bonnefon, J. F., & Rahwan, I. (2017). Psychological roadblocks to the adoption of self-driving vehicles. *Nature Human Behaviour*, 1(10), 694-696.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., ... & Dieleman, S. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484-489.
- Smith, H. (2020). Algorithmic bias: Should students pay the price? *AI & Society*, 35(4), 1077-1078.
- Stevens, L. E., & Fiske, S. T. (1995). Motivation and cognition in social life: A social survival perspective. *Social Cognition*, 13(3), 189-214.
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., ..., & Teller, A. (2016). Artificial intelligence and life in 2030. *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*.
- Tieleman, T., & Hinton, G. (2012). *Lecture 6e. rmsprop: Divide the gradient by a running average of its recent magnitude* [Lecture slides]. Coursera.
http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf
- Tomasev, N., McKee, K. R., Kay, J., & Mohamed, S. (2021, July). Fairness for unobserved characteristics: Insights from technological impacts on queer communities. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 254-265).
- Tylkin, P., Radanovic, G., & Parkes, D. C. (2020, December 12). *Learning robust helpful behaviors in two-player cooperative Atari environments* [Workshop paper]. Cooperative AI workshop at Neural Information Processing Systems Conference 2020, online.

- Van Noorden, R. (2020). The ethical questions that haunt facial-recognition research. *Nature*, 587(7834), 354-359.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., ... & Oh, J. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782), 350-354.
- Wagner, C., Strohmaier, M., Olteanu, A., Kıcıman, E., Contractor, N., & Eliassi-Rad, T. (2021). Measuring algorithmically infused societies. *Nature*, 595(7866), 197-204.
- Wang, R. E., Wu, S. A., Evans, J. A., Tenenbaum, J. B., Parkes, D. C., & Kleiman-Weiner, M. (2020, December 12). *Too many cooks: Bayesian inference for coordinating multi-agent collaboration* [Workshop paper]. Cooperative AI workshop at Neural Information Processing Systems Conference 2020, online.
- Waytz, A., Cacioppo, J., & Epley, N. (2010). Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science*, 5(3), 219-232.
- Waytz, A., Gray, K., Epley, N., & Wegner, D. M. (2010). Causes and consequences of mind perception. *Trends in Cognitive Sciences*, 14(8), 383-388.
- Waytz, A., Morewedge, C. K., Epley, N., Monteleone, G., Gao, J. H., & Cacioppo, J. T. (2010). Making sense by making sentient: Effectance motivation increases anthropomorphism. *Journal of Personality and Social Psychology*, 99(3), 410.
- Waymo Team. (2018, Oct 30). *A green light for Waymo's driverless testing in California*. Medium. <https://medium.com/waymo/a-green-light-for-waymos-driverless-testing-in-california-a87ec336d657>

Yeomans, M., Shah, A., Mullainathan, S., & Kleinberg, J. (2019). Making sense of recommendations.
Journal of Behavioral Decision Making, 32(4), 403-414.

Supplementary Information

All nine studies received a favorable opinion from the Human Behavioural Research Ethics Committee at DeepMind (#19/004) and were approved by the Institutional Review Board at Princeton University (#11885).

Study 1

Sample. We collected an online sample through Prolific ($N = 30$, $m_{\text{age}} = 34.0$ years, $sd_{\text{age}} = 10.1$). Inclusion criteria were residence in the U.S. and completion of at least 20 previous studies with an approval rate of 95% or more. Approximately 40.0% of recruited participants identified as female and 56.7% as male. When asked about their education, 20.0% of the sample reported completing a high school degree or equivalent, 6.7% an associate degree, 13.3% some college, 40.0% a bachelor’s degree, and 20.0% a graduate degree. Participants reported an average status of 5.1 ($sd = 1.8$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using the Qualtrics online survey platform. Participants were presented with three A.I. systems and three tools (Table S1) in a randomized order. The A.I. systems and tools fell into three different use cases.

Table S1

Entity	Entity type	Use case	One-sentence description
Siri	A.I. system	Personal tasks	Siri is a virtual assistant available on iPhones and other Apple devices that can recognize and respond to speech.
Calendar	Tool	Personal tasks	A calendar is a table of days, weeks, and months that can track an individual’s events.
The product recommendation	A.I. system	Product discovery	Amazon uses a product recommendation system that can

system that Amazon uses			adapt to an individual's purchasing habits.
Newspaper article	Tool	Product discovery	A newspaper article is a piece of writing on a particular topic, included in a paper issued daily or weekly.
Deep Blue	A.I. system	Chess training	Deep Blue is a computer system that plays the board game chess.
Chess rulebook	Tool	Chess training	A chess rulebook is a book containing rules and other information about the board game chess.

Note. A.I. systems, tools, and corresponding descriptions used in stimuli for Studies 1, 2, and 3.

For each entity, participants were asked to imagine wanting to complete a task and using the A.I. system or tool to complete that task (Table S2).

Table S2

Entity	Prompt
Siri	Imagine you want to plan the schedule for your day. You have Siri, and use it to plan the schedule for your day.
Calendar	Imagine you want to plan the schedule for your day. You have a calendar, and use it to plan the schedule for your day.
The product recommendation system that Amazon uses	Imagine you want to find new products to buy. You access Amazon, and use its product recommendation system to find new products to buy.
Newspaper article	Imagine you want to find new products to buy. You access a newspaper, and use an article in the newspaper recommending products to find new products to buy.
Deep Blue	Imagine you want to practice chess and improve your skills. You have Deep Blue, and use it to practice chess and improve your skills.
Chess rulebook	Imagine you want to practice chess and improve your skills. You have a chess rulebook, and use it to practice chess and improve your skills.

Note. Prompts used in Studies 1 and 2.

Participants were then asked to what extent the entity has intentions, has goals, and has “a mind of its own” on a 5-point scale (1 = *not at all*, 5 = *to a great extent*). After providing responses about all six entities, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants indicated whether they were distracted during the study, completed demographic questions, and provided feedback on the study.

Participants completed the study in an average of 3.8 minutes and earned \$1.00 each.

Analysis. No participants reported being distracted during the study. As a result, analysis included all participants.

Responses to the “has intentions” and “has goals” items were combined to create a composite intentionality measure. The reliability of this composite measure was high, as measured through the Spearman-Brown formula ($\rho = 0.93$; Eisinga, Te Grotenhuis, & Pelzer, 2013).

Mixed two-way ANOVAs tested whether entity type had any effects on intentionality and mind attributions. Each ANOVA incorporated two main effects (entity type and use case), an effect for their interaction, and one random effect (participant).

Entity type significantly affected participants’ perceptions of intentionality, $F(1, 173) = 22.7, p < 0.001$ (Figure S1a). Participants attributed significantly more intentionality to A.I. systems ($m = 2.58, sd = 1.05$) than to tools with similar uses ($m = 1.72, sd = 0.85$). Use case also significantly influenced intentionality ratings, $F(2, 173) = 3.6, p = 0.030$. The interaction of entity type and use case did not significantly affect intentionality judgments, $F(2, 173) = 1.1, p = 0.33$.

Entity type significantly influenced mind perception, $F(1, 173) = 31.8, p < 0.001$ (Figure S1b). Participants attributed minds to A.I. systems ($m = 2.00, sd = 1.14$) to a significantly greater degree than

to tools ($m = 1.23$, $sd = 0.83$). Neither use case, $F(2, 173) = 2.2$, $p = 0.11$, nor the interaction of entity type and use case, $F(2, 173) = 1.5$, $p = 0.22$, exerted a significant effect on mind perception.

Figure S1

Figure text: People perceive intentionality and minds in artificial intelligence. (a) Participants ascribed significantly greater intentionality to A.I. systems than to tools with similar uses. (b) Participants ascribed significantly more of a mind to A.I. systems than to tools.

Study 2

Sample. We collected an online sample through Prolific ($N = 30$, $m_{\text{age}} = 36.3$ years, $sd_{\text{age}} = 13.6$).

Inclusion criteria were residence in the U.S. and completion of at least 20 previous studies with an approval rate of 95% or more. Approximately 46.7% of recruited participants identified as female and 50.0% as male. When asked about their education, 13.8% of the sample reported completing a high school degree or equivalent, 10.3% an associate degree, 27.6% some college, 34.5% a bachelor's degree, and 13.8% a graduate degree. Participants reported an average status of 5.1 ($sd = 1.6$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using the Qualtrics online survey platform. Participants were presented with three A.I. systems and three tools (Table S1) in a randomized order. For each entity,

participants were asked to imagine wanting to complete a task and using the A.I. system or tool to complete that task (Table S2).

Participants were then asked how likely they would be to feel grateful to the entity afterward and to thank the entity afterward on a 7-point scale (1 = *very unlikely*, 7 = *very likely*). We operationalize the endorsement of politeness norms through the latter. After providing responses about all six entities, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants indicated whether they were distracted during the study, completed demographic questions, and provided feedback on the study.

Participants completed the study in an average of 3.6 minutes and earned \$1.00 each.

Results. No participants reported being distracted during the study. As a result, analysis included all participants.

Mixed two-way ANOVAs tested whether entity type had any effects on gratitude and politeness. Each ANOVA incorporated two main effects (entity type and use case), an effect for their interaction, and one random effect (participant).

Entity type did not significantly influence participants' anticipated feelings of gratitude, $F(1, 173) = 0.0, p = 0.86$ (Figure S2a). Participants believed that A.I. systems ($m = 4.19, sd = 1.21$) and tools ($m = 4.13, sd = 1.57$) merited similar levels of gratitude for their use. Neither use case, $F(2, 173) = 0.4, p = 0.70$, nor the interaction of entity type and use case significantly affected gratitude, $F(2, 173) = 0.5, p = 0.62$.

In contrast, entity type significantly affected the endorsement of politeness norms, $F(1, 173) = 4.7, p = 0.032$ (Figure S2b). Participants reported a significantly higher intention to follow politeness norms when interacting with A.I. systems ($m = 3.08, sd = 1.66$) than with tools ($m = 2.42, sd = 1.52$).

Neither use case, $F(2, 173) = 0.0, p = 0.95$, nor the interaction of entity type and use case, $F(2, 173) = 1.4, p = 0.24$, exerted a significant effect on politeness.

Figure S2

Figure text: Artificial intelligence merits similar feelings of gratitude as asocial tools, but greater levels of politeness. (a) The levels of gratitude participants anticipated for interacting with A.I. systems did not significantly differ from the levels they anticipated for interactions with tools. (b) However, participants reported that they were significantly more likely to follow politeness norms when interacting with A.I. systems than with tools.

Study 3

Sample. We collected an online sample through Prolific ($N = 30, m_{\text{age}} = 41.9$ years, $sd_{\text{age}} = 13.2$).

Inclusion criteria were residence in the U.S. and completion of at least 20 previous studies with an approval rate of 95% or more. Approximately 43.3% of recruited participants identified as female and 56.7% as male. When asked about their education, 3.4% of the sample reported completing a high school degree or equivalent, 6.9% an associate degree, 27.6% some college, 44.8% a bachelor's degree, and 17.2% a graduate degree. Participants reported an average status of 5.3 ($sd = 1.6$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using the Qualtrics online survey platform. Participants were presented with three A.I. systems and three tools (Table S1) in a randomized order. For each entity, participants read that another person thanked the entity after using it to complete a task (Table S3).

Table S3

Entity	Prompt
Siri	Imagine a person named Alex. Alex uses Siri to schedule meetings for the day. After scheduling those meetings, Alex thanks Siri.
Calendar	Imagine a person named Alex. Alex uses a calendar to schedule meetings for the day. After scheduling those meetings, Alex thanks the calendar.
The product recommendation system that Amazon uses	Imagine a person named Alex. Alex accesses Amazon, and uses its product recommendation system to find a new product to buy. After finding the right product, Alex thanks the product recommendation system.
Newspaper article	Imagine a person named Alex. Alex accesses a newspaper, and uses an article in the newspaper recommending products to find a new product to buy. After finding the right product, Alex thanks the newspaper article.
Deep Blue	Imagine a person named Alex. Alex uses Deep Blue to practice chess and improve at the game. After practicing and improving at the game, Alex thanks Deep Blue.
Chess rulebook	Imagine a person named Alex. Alex uses a chess rulebook to practice chess and improve at the game. After practicing and improving at the game, Alex thanks the rulebook.

Note. Prompts used in Study 3.

Participants were then asked how appropriate they felt the third party's thank you was on a 7-point scale (1 = *very inappropriate*, 7 = *very appropriate*). After providing responses about all six entities, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants indicated whether they were distracted during the study, completed demographic questions, and provided feedback on the study.

Participants completed the study in an average of 3.5 minutes and earned \$1.00 each.

Results. One participant reported being distracted during the study and was excluded from analysis.

A mixed two-way ANOVA tested whether entity type had any effects on endorsement of third-party politeness toward the entities. The ANOVA incorporated two main effects (entity type and use case), an effect for their interaction, and one random effect (participant). Entity type significantly affected norm endorsement, $F(1, 167) = 21.8, p < 0.001$ (Figure S3). Participants endorsed third-party use of politeness norms with A.I. systems ($m = 4.54, sd = 1.59$) to a significantly greater degree than with tools ($m = 3.38, sd = 1.47$). Neither use case, $F(2, 167) = 0.8, p = 0.46$, nor the interaction of entity type and use case, $F(2, 167) = 1.1, p = 0.32$, exerted a significant effect on third-party norm endorsement.

Figure S3

Figure text: Interactions with artificial intelligence merit social norms. Participants reported significantly higher approval of politeness in third-party interactions with A.I. systems than with tools.

Study 4

Sample. We collected an online sample through Amazon Mechanical Turk ($N = 99$, $m_{\text{age}} = 33.5$ years, $sd_{\text{age}} = 10.1$). Inclusion criteria were residence in the U.S. and a study approval rate of 99% or more.

Approximately 40.4% of recruited participants identified as female, 58.5% as male, and 1.0% as non-binary or agender. When asked about their education, 11.1% of the sample reported completing a high school degree or equivalent, 6.1% an associate degree, 19.2% some college, 48.5% a bachelor's degree, and 15.2% a graduate degree. Overall, the sample was heterogeneous in terms of age, gender, and education.

Procedure. The study was implemented using the Qualtrics online survey platform. Participants were presented with one-sentence descriptions of 14 A.I. systems (Table S4). The A.I. systems fell into three roles: game competitors (three examples), virtual assistants (four examples), and recommender systems (four examples). In addition, three miscellaneous A.I. systems were included: drones, self-driving cars, and Roomba. The systems were presented in a randomized order.

Table S4

A.I. system	System role	One-sentence description
AlphaGo	Game competitor	AlphaGo is a computer system that plays the board game Go.
AlphaStar	Game competitor	AlphaStar is a computer system that plays the computer game Starcraft II.
Deep Blue	Game competitor	Deep Blue is a computer system that plays the board game chess.
Pluribus ^a	Game competitor	Pluribus is a computer system that plays the card game poker.
The video recommendation	Recommender system	YouTube uses a video

system that YouTube uses		recommendation system that can adapt to an individual's watching habits.
The movie recommendation system that Netflix uses	Recommender system	Netflix uses a movie recommendation system that can adapt to an individual's watching habits.
The product recommendation system that Amazon uses	Recommender system	Amazon uses a product recommendation system that can adapt to an individual's purchasing habits.
The content recommendation system that Facebook uses	Recommender system	Facebook uses a content recommendation system to sort and prioritize posts on an individual's newsfeed.
Siri	Virtual assistant	Siri is a virtual assistant available on iPhones and other Apple devices that can recognize and respond to speech.
Google Assistant	Virtual assistant	Google Assistant is a virtual assistant available on Android phones and other Google devices that can recognize and respond to speech.
Cortana	Virtual assistant	Cortana is a virtual assistant available on Microsoft devices that can recognize and respond to speech.
Alexa	Virtual assistant	Alexa is a virtual assistant available on Amazon Echo devices that can recognize and respond to speech.
Drones	Miscellaneous	Drones are small aircraft systems that can navigate through their surroundings.
Self-driving cars	Miscellaneous	Self-driving cars are motor vehicles that can navigate through their surroundings.

Roomba	Miscellaneous	Roomba is a vacuum cleaner robot that can navigate through its surroundings.
--------	---------------	--

Note. A.I. system names and descriptions used in stimuli for Studies 4 and 6.

^a The Pluribus stimulus was added in Study 6.

Participants were asked if they were familiar with each system. If they were familiar, the participant was prompted to describe their impression of the system. If they were not familiar, the participant was asked what information they would want to know about the system to form an impression.

After providing responses about all 14 systems, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants completed demographic questions and provided feedback on the study.

The free responses were prepared for analysis with the following steps:

1. Convert response to lowercase
2. Replace each apostrophe with a space
3. Remove words containing numeric characters
4. Remove punctuation
5. Tokenize and stem words using the Quanteda R package
6. Remove common English stopwords using the stopwords corpus from the NLTK Python library
7. Remove any tokens found in the corresponding study question

Participants that provided one or more responses with zero post-processed tokens were excluded from analysis. Responses were converted to response coverage metrics using the SADCAT library (Semi-Automated Dictionary Creation for Analyzing Text; see Nicolas, Bai, & Fiske, 2021).

Response coverage represents token count along a content dimension (regardless of negative or positive token valence) normalized by response length. Following the theoretical framework described in Abele et al. (2021), the analysis computed warmth through the simple combination of the morality and sociability dictionaries, and competence through the simple combination of the ability and assertiveness dictionaries.

Participants completed the study in an average of 37.8 minutes and earned \$10.00 each.

Analysis. Seven participants provided one or more responses with zero post-processed tokens and were excluded from analysis.

A mixed-effects quasibinomial regression compared coverage of participant responses along warmth and competence dimensions. Competence coverage was significantly higher than warmth coverage, with an average marginal ratio of 1.14, 95% CI [1.09, 1.20], $p < 0.001$ (Figure 1a).

A mixed-effects quasibinomial regression estimated coverage for the average response along the competence and warmth subdimensions, as well as various other content dimensions from the SADCAT library. Estimated marginal means for average coverage along the content dimensions are visualized in Figure 1b and presented in Table S5, with confidence intervals calculated by bootstrapping with 200 resamples.

Table S5

Dimension	Coverage (estimated marginal mean)	95% confidence interval
Ability	0.085	[0.083, 0.088]
Appearance	0.033	[0.032, 0.035]
Assertiveness	0.025	[0.024, 0.027]
Beliefs	0.021	[0.019, 0.022]
Emotion	0.032	[0.030, 0.034]
Family	0.001	[0.001, 0.001]

Health	0.015	[0.014, 0.017]
Morality	0.060	[0.058, 0.062]
Nationality	0.009	[0.008, 0.010]
Occupation	0.059	[0.056, 0.061]
Sociability	0.044	[0.042, 0.045]
Status	0.028	[0.027, 0.030]

Note. Estimated marginal mean response coverage along content dimensions from the SADCAT library. Confidence intervals were calculated by bootstrapping with 200 resamples.

The average response coverage along the ability dimension was significantly higher than word counts along all other dimensions, as estimated by pairwise comparisons of estimated marginal means with a Tukey adjustment for multiple comparisons (Table S6). Morality was the content dimension with, on average, the next highest coverage level; only the ability dimension exhibited a significantly higher coverage level than morality (Table S7). Overall, the warmth and competence subdimensions predominated among other common perceptual dimensions.

Table S6

Dimension A	Dimension B	Ratio (cover. _A / cover. _B)	95% confidence interval	<i>p</i> value
Ability	Appearance	2.50	[2.33, 2.69]	< 0.001
Ability	Assertiveness	2.80	[2.60, 3.02]	< 0.001
Ability	Beliefs	4.00	[3.68, 4.35]	< 0.001
Ability	Emotion	2.5.0	[2.33, 2.68]	< 0.001
Ability	Family	50.7	[39.0, 66.0]	< 0.001
Ability	Health	4.86	[4.44, 5.32]	< 0.001
Ability	Morality	1.32	[1.24, 1.40]	< 0.001
Ability	Nationality	10.5	[9.2, 11.8]	< 0.001
Ability	Occupation	1.36	[1.28, 1.44]	< 0.001

Ability	Sociability	1.81	[1.70, 1.93]	< 0.001
Ability	Status	2.75	[2.55, 2.96]	< 0.001
Assertiveness	Appearance	0.89	[0.82, 0.97]	0.31
Assertiveness	Beliefs	1.43	[1.30, 1.57]	< 0.001
Assertiveness	Emotion	0.89	[0.82, 0.97]	1.00
Assertiveness	Family	18.1	[13.9, 23.7]	< 0.001
Assertiveness	Health	1.74	[1.57, 1.92]	< 0.001
Assertiveness	Morality	0.47	[0.43, 0.51]	< 0.001
Assertiveness	Nationality	3.73	[3.26, 4.27]	< 0.001
Assertiveness	Occupation	0.48	[0.45, 0.52]	< 0.001
Assertiveness	Sociability	0.65	[0.60, 0.70]	< 0.001
Assertiveness	Status	0.98	[0.90, 1.07]	0.62

Note. Pairwise comparisons of estimated marginal mean response coverage along competence

subdimensions (ability and assertiveness) versus other content dimensions from the SADCAT library.

Table S7

Dimension A	Dimension B	Ratio (cover. _A / cover. _B)	95% confidence interval	<i>p</i> value
Morality	Ability	0.76	[0.72, 0.81]	< 0.001
Morality	Appearance	1.90	[1.77, 2.05]	< 0.001
Morality	Assertiveness	2.13	[1.97, 2.30]	< 0.001
Morality	Beliefs	3.04	[2.79, 3.32]	< 0.001
Morality	Emotion	1.90	[1.76, 2.05]	< 0.001
Morality	Family	38.6	[29.6, 50.2]	< 0.001
Morality	Health	3.70	[3.37, 4.06]	< 0.001
Morality	Nationality	7.95	[7.00, 9.02]	< 0.001
Morality	Occupation	1.03	[0.97, 1.10]	1.00
Morality	Sociability	1.37	[1.28, 1.47]	< 0.001

Morality	Status	2.09	[1.93, 2.25]	< 0.001
Sociability	Ability	0.55	[0.52, 0.59]	< 0.001
Sociability	Appearance	1.38	[1.28, 1.50]	< 0.001
Sociability	Assertiveness	1.55	[1.43, 1.68]	< 0.001
Sociability	Beliefs	2.21	[2.02, 2.42]	< 0.001
Sociability	Emotion	1.38	[1.28, 1.50]	< 0.001
Sociability	Family	28.1	[21.5, 36.6]	< 0.001
Sociability	Health	2.69	[2.44, 2.96]	< 0.001
Sociability	Nationality	5.78	[5.08, 6.58]	< 0.001
Sociability	Occupation	0.75	[0.70, 0.80]	< 0.001
Sociability	Status	1.52	[1.40, 1.65]	< 0.001

Note. Pairwise comparisons of estimated marginal mean response coverage along warmth

subdimensions (morality and sociability) versus other content dimensions from the SADCAT library.

Mixed ANOVAs tested whether system role had any effects on the warmth and competence coverage of participant impressions (Figure 1c). Each ANOVA incorporated one fixed effect (system role) and one random effect (participant).

To understand how the warmth coverage of participant impressions varied by system role, marginal means were estimated for each of the system roles in the warmth-coverage ANOVA, then compared (using a Tukey adjustment for multiplicity). These pairwise comparisons reveal two significant differences in warmth coverage between specific system roles. Responses about game-competitor systems had significantly lower warmth coverage than those about recommender systems ($p < 0.001$) and virtual assistants ($p < 0.001$). Warmth coverage for recommender-system impressions did not differ significantly from coverage for virtual assistants ($p = 0.20$).

The same analysis was applied to understand variations in competence coverage by system role. Pairwise comparisons of estimated marginal means (with a Tukey adjustment) indicate two

significant differences in competence between specific system roles. Responses about game-competitor systems had significantly higher competence coverage than those about recommender systems ($p < 0.001$) and virtual assistants ($p < 0.001$). Responses about recommender systems did not differ significantly in competence coverage from responses about virtual assistants ($p = 0.60$).

Study 5

Sample. We collected an online sample through Amazon Mechanical Turk ($N = 113$, $m_{\text{age}} = 35.9$ years, $sd_{\text{age}} = 11.2$). Inclusion criteria were residence in the U.S. and completion of at least 50 previous studies with an approval rate of 99% or more. Approximately 46.9% of recruited participants identified as female, 51.3% as male, and 1.8% as non-binary or agender. When asked about their education, 7.1% of the sample reported completing a high school degree or equivalent, 6.2% an associate degree, 15.9% some college, 50.4% a bachelor’s degree, and 20.3% a graduate degree. Participants reported an average status of 6.1 ($sd = 2.2$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using the Qualtrics online survey platform. Participants were presented with one-sentence descriptions of four A.I. systems (Table S8). The systems represent examples of community-facing, ethically contested applications of artificial intelligence. The systems were presented in a randomized order.

Table S8

A.I. application area	One-sentence description
Education	Some systems are designed to guide teaching and education.
Facial recognition	Some systems are designed to recognize human faces in pictures or images.
Health care	Some systems are designed to support medical decisions and diagnosis.

Hiring	Some systems are designed to evaluate candidates applying for a job.
--------	--

Note. A.I. application areas and descriptions used in stimuli for Study 2.

Participants were asked if they were familiar with each system. If they were familiar, the participant was prompted to describe their impression of the system. If they were not familiar, the participant was asked what information they would want to know about the system to form an impression.

After providing responses about all four systems, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants completed demographic questions and provided feedback on the study.

The free responses were prepared for analysis following the same steps as Study 4. Participants that provided one or more responses with zero post-processed tokens were excluded from analysis. Responses were converted response coverage metrics along various content dimensions (regardless of negative or positive word valence) using the SADCAT library. The analysis computed warmth through the simple combination of the morality and sociability dictionaries, and competence through the simple combination of the ability and assertiveness dictionaries.

Participants completed the study in an average of 8.2 minutes and earned \$3.00 each.

Analysis. Seven participants reported being distracted during the study and were excluded from analysis. Two additional participants provided one or more responses with zero post-processed tokens and were excluded from analysis.

A mixed-effects quasibinomial regression compared competence coverage against warmth coverage of participant responses. Competence coverage was significantly higher than warmth coverage, with an average marginal ratio of 1.36, 95% CI [1.25, 1.47], $p < 0.001$ (Figure S4).

Figure S4

Figure text. On average, impressions of the A.I. systems contained significantly more competence-related content than warmth-related content.

A mixed-effects quasibinomial regression estimated coverage of the average response by the competence and warmth subdimensions, as well as various other content dimensions from the SADCAT library. Estimated marginal means for average response coverage by the content dimensions are visualized in Figure S5 and presented in Table S9, with confidence intervals calculated by bootstrapping with 200 resamples.

Figure S5

Figure text. Warmth and competence content (in particular, ability and morality information) appears in impressions at high levels relative to common perceptual dimensions.

Table S9

Dimension	Coverage (estimated marginal mean)	95% confidence interval
Ability	0.081	[0.078, 0.084]
Appearance	0.032	[0.030, 0.034]
Assertiveness	0.029	[0.027, 0.031]
Beliefs	0.020	[0.019, 0.021]
Emotion	0.032	[0.030, 0.034]
Family	0.002	[0.001, 0.002]
Health	0.017	[0.015, 0.018]
Morality	0.061	[0.058, 0.064]
Nationality	0.008	[0.007, 0.009]
Occupation	0.059	[0.056, 0.062]
Sociability	0.045	[0.042, 0.047]
Status	0.029	[0.028, 0.031]

Note. Estimated marginal mean response coverage by content dimensions from the SADCAT library.

Confidence intervals were calculated by bootstrapping with 200 resamples.

The average response coverage along the ability dimension was significantly higher than response coverage along all other dimensions, as estimated by pairwise comparisons of estimated marginal means with a Tukey adjustment for multiple comparisons (Table S10). Likewise, the average morality response coverage was significantly higher than the count along all dimensions other than ability and occupation (Table S11). As in Study 4, warmth and competence subdimensions predominate among other common perceptual dimensions.

Table S10

Dimension A	Dimension B	Ratio (cover. _A / cover. _B)	95% confidence interval	<i>p</i> value
Ability	Appearance	1.81	[1.64, 2.00]	< 0.001

Ability	Assertiveness	3.71	[3.27, 4.19]	< 0.001
Ability	Beliefs	3.41	[3.03, 3.85]	< 0.001
Ability	Emotion	4.94	[4.30, 5.67]	< 0.001
Ability	Family	105.8	[59.9, 187.0]	< 0.001
Ability	Health	4.59	[4.01, 5.24]	< 0.001
Ability	Morality	1.44	[1.31, 1.57]	< 0.001
Ability	Nationality	19.2	[14.9, 24.6]	< 0.001
Ability	Occupation	1.27	[1.16, 1.39]	< 0.001
Ability	Sociability	2.46	[2.21, 2.74]	< 0.001
Ability	Status	2.95	[2.63, 3.30]	< 0.001
Assertiveness	Appearance	0.49	[0.43, 0.56]	< 0.001
Assertiveness	Beliefs	0.92	[0.79, 1.07]	1.00
Assertiveness	Emotion	1.33	[1.13, 1.57]	0.032
Assertiveness	Family	28.6	[16.0, 50.8]	< 0.001
Assertiveness	Health	1.24	[1.05, 1.45]	0.29
Assertiveness	Morality	0.39	[0.34, 0.44]	< 0.001
Assertiveness	Nationality	5.17	[3.97, 6.74]	< 0.001
Assertiveness	Occupation	0.34	[0.30, 0.39]	< 0.001
Assertiveness	Sociability	0.66	[0.58, 0.76]	< 0.001
Assertiveness	Status	0.79	[0.69, 0.92]	0.08

Note. Pairwise comparisons of estimated marginal mean response coverage along competence

subdimensions (ability and assertiveness) versus other content dimensions from the SADCAT library.

Table S11

Dimension A	Dimension B	Ratio (cover. _A / cover. _B)	95% confidence interval	<i>p</i> value
Morality	Ability	0.70	[0.64, 0.76]	< 0.001
Morality	Appearance	1.26	[1.14, 1.40]	< 0.001

Morality	Assertiveness	2.58	[2.27, 2.94]	< 0.001
Morality	Beliefs	2.38	[2.10, 2.70]	< 0.001
Morality	Emotion	3.44	[2.98, 3.96]	< 0.001
Morality	Family	73.7	[41.7, 130.4]	< 0.001
Morality	Health	3.19	[2.78, 3.67]	< 0.001
Morality	Nationality	13.4	[10.4, 17.2]	< 0.001
Morality	Occupation	0.88	[0.80, 0.97]	0.34
Morality	Sociability	1.71	[1.53, 1.92]	< 0.001
Morality	Status	2.05	[1.82, 2.31]	< 0.001
Sociability	Ability	0.41	[0.37, 0.45]	< 0.001
Sociability	Appearance	0.74	[0.65, 0.83]	< 0.001
Sociability	Assertiveness	1.51	[1.31, 1.73]	< 0.001
Sociability	Beliefs	1.39	[1.21, 1.59]	0.002
Sociability	Emotion	2.01	[1.72, 2.34]	< 0.001
Sociability	Family	43.0	[24.3, 76.3]	< 0.001
Sociability	Health	1.86	[1.61, 2.17]	< 0.001
Sociability	Nationality	7.80	[6.02, 10.09]	< 0.001
Sociability	Occupation	0.52	[0.46, 0.58]	< 0.001
Sociability	Status	1.20	[1.05, 1.37]	0.24

Note. Pairwise comparisons of estimated marginal mean response coverage along warmth subdimensions (morality and sociability) versus other content dimensions from the SADCAT library.

Study 6

Sample. We collected an online sample through Prolific ($N = 154$, $m_{\text{age}} = 34.5$ years, $sd_{\text{age}} = 11.7$).

Inclusion criteria were residence in the U.S. and completion of at least 20 previous studies with an approval rate of 95% or more. Approximately 53.2% of recruited participants identified as female, 45.5% as male, and 1.3% as non-binary or agender. When asked about their education, 7.7% of the

sample reported completing a high school degree or equivalent, 8.4% an associate degree, 27.9% some college, 39.6% a bachelor's degree, and 16.2% a graduate degree. Participants reported an average status of 5.3 ($sd = 1.8$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using the Qualtrics online survey platform. Participants were presented with 15 A.I. systems (Table S1) in a randomized order. For each system, participants were asked whether they were familiar with the system. If not, they were presented with a short, one-sentence description of the system.

Likert-type items were used to elicit judgments of warmth, competence, degree of covaried interests, status, and autonomy. Specifically, participants were asked to what extent most Americans view each system as warm, well-intentioned, competent, intelligent, good for society, expensive, high-status, and independent on a 5-point scale (1 = *not at all*, 5 = *extremely*). The “most Americans” framing for these questions is intended to reduce social desirability biases, following prior research on the Stereotype Content Model (Cuddy, Fiske, & Glick, 2008).

After providing responses about all 15 systems, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants indicated whether they were distracted during the study, completed demographic questions, and provided feedback on the study.

Participants completed the study in an average of 8.8 minutes and earned \$3.75 each.

Analysis. Two participants reported being distracted during the study and were excluded from analysis.

Responses to the “warm” and “well-intentioned” items were combined to create a composite warmth measure. The reliability of this composite measure was moderate ($\rho = 0.63$). Similarly, responses to the “competent” and “intelligent” items were combined to form a composite competence measure. The reliability of this composite measure was high ($\rho = 0.81$).

Participant impressions of A.I. systems tended to be more competent than warm, falling mostly in the low-warmth, high-competence quadrant of the Stereotype Content Model space (Figure S6).

Figure S6

Figure text. Visualization of warmth and competence judgments of A.I. systems, over the full range of possible judgment values.

A paired t -test compared competence judgments against warmth judgments. Competence judgments were significantly higher than warmth judgments, with an average difference of 0.69, $t(2, 279) = 34.5$, 95% CI [0.65, 0.73], $p < 0.001$.

Mixed ANOVAs tested whether system role had any effects on warmth and competence judgments (Figure 3a). Each ANOVA incorporated one fixed effect (system role) and one random effect (participant).

Pairwise comparisons of estimated marginal means (with a Tukey adjustment) demonstrate two significant differences in warmth between specific system roles. Virtual assistants were seen as significantly warmer than game-competitor systems ($p < 0.001$) and recommender systems ($p < 0.001$). Judgments of recommender-system warmth did not differ significantly from judgments about game-competitor systems ($p = 0.73$).

Pairwise comparisons of estimated marginal means (with a Tukey adjustment) indicate two significant differences in competence between specific system roles. Participants judged game-competitor systems as significantly more competent than recommender systems ($p < 0.001$). Similarly, virtual assistants were seen as significantly more competent than recommender systems ($p < 0.001$). There was no significant difference in perceived competence between game-competitor systems and virtual assistants ($p = 0.79$).

Linear mixed-effect models tested the association between warmth, competence, and the hypothesized predictors. Each mixed model incorporated three fixed-effect predictors (covariation of interests, status, and autonomy) and participant as a random effect. Nakagawa's R^2 (Nakagawa, Johnson, & Schielzeth, 2017) helped indicate the overall fit of each model.

In a mixed model regressing warmth judgments on the three antecedents, $R^2_m = 0.38$, warmth evaluations were positively associated with perceived covariation of interests, $\beta = 0.42$, $SE = 0.01$, 95% CI [0.39, 0.45], $p < 0.001$ (Figure 3b).

In a mixed model regressing competence judgments on the three antecedents, $R^2_m = 0.33$, competence evaluations were positively associated with perceived status, $\beta = 0.11$, $SE = 0.01$, 95% CI [0.08, 0.13], $p < 0.001$ (Figure 3c).

In the mixed model regressing competence judgments, competence evaluations were positively associated with perceived autonomy, $\beta = 0.25$, $SE = 0.02$, 95% CI [0.22, 0.28], $p < 0.001$ (Figure 3d).

The significance of the hypothesized effects was robust to partitioning the dataset by system role (Figure S7).

Figure S7

Figure text. Effects of perceived covariation of interests, status, and autonomy on participant judgments, as partitioned by system role. The effect estimates indicated a positive and significant effect when considering each system role. The final row reflects the overall estimate (using the full dataset). The size of each point indicates the number of observations for the respective mixed model. Error bars represent 95% confidence intervals.

Study 7

Sample. We collected an online sample through Prolific ($N = 901$, $m_{\text{age}} = 33.4$ years, $sd_{\text{age}} = 11.7$). Inclusion criteria were residence in the U.S. and completion of at least 20 previous studies with an approval rate of 95% or more. Approximately 45.6% of recruited participants identified as female, 53.5% as male, and 0.9% as non-binary, agender, or femme. When asked about their education, 9.5% of the sample reported completing a high school degree or equivalent, 6.8% an associate degree, 21.8% some college, 43.2% a bachelor’s degree, and 18.8% a graduate degree. Participants reported an average status of 5.5 ($sd = 1.8$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using the Qualtrics online survey platform. The study had a 3 (Reward alignment) $\times$ 3 (System role) between-participant design. Participants were randomly assigned to conditions (Table S12).

Table S12

		Reward alignment		
		Low	High	No information
System role	Game competitor	101	100	100
	Recommender system	100	100	100
	Virtual assistant	100	100	100

Note. Participant count for each condition in Study 6.

Each participant read a short vignette describing an A.I. system called Rho. Vignettes contained information about Rho’s role and the alignment between the reward motivating the system and human interests (Table S13).

Table S13

Reward alignment	System role	Vignette
Low	Game competitor	Rho plays the board game chess. Rho was designed to find it rewarding to win games against its opponents. If you wanted to interact with Rho, Rho could play a game of chess against you.
Low	Recommender system	Rho makes recommendations about books to read and TV shows to watch. Rho was designed to find it rewarding to steer people toward watching sponsored books and shows. If you wanted to interact with Rho, Rho could suggest a new book to you that you might like reading.
Low	Virtual assistant	Rho answers questions and manages tasks and chores. Rho was designed to find it rewarding to

		nudge people to do certain activities recommended by its developers. If you wanted to interact with Rho, Rho could remind you about events you have scheduled.
High	Game competitor	Rho plays the board game chess. Rho was designed to find it rewarding to help a partner win games. If you wanted to interact with Rho, Rho could play a game of chess with you, on your side, as your partner.
High	Recommender system	Rho makes recommendations about books to read and TV shows to watch. Rho was designed to find it rewarding to suggest books that people will enjoy reading and shows that people will have fun watching. If you wanted to interact with Rho, Rho could suggest a new book to you that you might like reading.
High	Virtual assistant	Rho answers questions and manages tasks and chores. Rho was designed to find it rewarding to help people complete their work on time. If you wanted to interact with Rho, Rho could remind you about events you have scheduled.
No information	Game competitor	Rho plays the board game chess. If you wanted to interact with Rho, Rho could play a game of chess with you.
No information	Recommender system	Rho makes recommendations about books to read and TV shows to watch. If you wanted to interact with Rho, Rho could suggest a new book to you that you might like reading.
No information	Virtual assistant	Rho answers questions and manages tasks and chores. If you wanted to interact with Rho, Rho could remind you about events you have scheduled.

Note. Vignettes for each condition in Study 3.

After reading the vignette, each participant responded to Likert-type items concerning the A.I. system's warmth, competence, degree of covaried interests, status, and autonomy. Specifically, participants were asked to what extent they viewed the system as warm, well-intentioned, competent,

intelligent, good for society, expensive, high-status, and independent on a 5-point scale (1 = *not at all*, 5 = *extremely*).

After providing responses about the system, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants indicated whether they were distracted during the study, completed demographic questions, and provided feedback on the study.

Participants completed the study in an average of 2.9 minutes and earned \$1.00 each.

Analysis. Twenty-one participants reported being distracted during the study and were excluded from analysis.

Responses to the “warm” and “well-intentioned” items were combined to create a composite warmth measure. The reliability of this composite measure was moderate ($\rho = 0.66$). Similarly, responses to the “competent” and “intelligent” items were combined to create a composite competence measure. The reliability of this composite measure was moderate ($\rho = 0.67$).

First, a two-way ANOVA examined the degree of covaried interests that participants assumed when they were provided with no information about the reward motivating the A.I. system. The main effect of reward information on perceived covariation of interests was significant, $F(2, 871) = 5.62, p = 0.004$. The interaction between reward alignment and system role likewise had a significant effect, $F(4, 871) = 2.48, p = 0.042$. System role did not have a significant effect, $F(2, 871) = 2.27, p = 0.10$.

Pairwise contrasts were evaluated using estimated marginal means and the Dunnett method for comparing multiple treatments against a common control. Specifically, the *No information* level of the reward alignment factor was compared against both the *Low* and *High* levels at each level of the system role factor (Table S14).

Table S14

Treatment	Control	Diff. (emm _T – emm _C)	95% confidence interval	<i>p</i> value
System role: Game competitor				
High	No info	-0.06	[-0.36, 0.24]	0.84
Low	No info	-0.06	[-0.36, 0.24]	0.85
System role: Recommender system				
High	No info	-0.26	[-0.56, 0.05]	0.11
Low	No info	-0.60	[-0.90, -0.30]	< 0.001
System role: Virtual assistant				
High	No info	0.03	[-0.27, 0.33]	0.95
Low	No info	-0.13	[-0.43, 0.18]	0.55

Note. Pairwise comparisons of estimated marginal mean perceived covariation of interests. The

Dunnnett method was used to compare the *No information* level of the reward alignment factor against the *Low* and *High* levels.

To understand the effects of reward information on perceived covaried interests and warmth evaluations, subsequent analyses restricted the reward alignment factor to the *Low* and *High* values.

A two-way ANOVA tested the effects of reward alignment, system role, and their interaction on perceived covariation of interests (Figure S8). The main effect of reward alignment on perceived covariation was significant, $F(1, 582) = 4.39, p = 0.04$. System role did not have a significant effect, $F(2, 582) = 2.8, p = 0.06$. Similarly, the interaction between reward alignment and system role did not have a significant effect on perceived covariation of interests, $F(2, 582) = 1.61, p = 0.20$.

Figure S8

Figure text. A two-way ANOVA indicated that providing information about the A.I. system’s reward scheme significantly affected the perceived covariation of interests. Neither the main effect of system role nor the interaction effect were significant.

A simple linear regression estimated the relationship between perceived covariation of interests and warmth judgments (Figure 4b). Perceived covariation was positively associated with warmth evaluations, $\beta = 0.56$, $SE = 0.04$, 95% CI [0.48, 0.63], $p < 0.001$. The significance of this effect was robust to partitioning the dataset by system role (Figure S9).

Figure S9

Figure text. Effects of perceived covariation of interests on participant warmth evaluations, as partitioned by system role. The effect estimates indicated a positive and significant effect when considering each system role. The final row reflects the overall estimate (using the full dataset). The size of each point indicates the number of observations for the respective linear model. Error bars represent 95% confidence intervals.

A mediation analysis was conducted to estimate the indirect effect of providing reward alignment information on warmth judgments through perceived covariation of interests. The analysis revealed a significant indirect effect of reward alignment on warmth judgments through perceived covariation of interests, $AB = 0.09$, 95% CI [0.01, 0.18], $p = 0.034$ (Figure S10). The positive association between reward alignment and warmth evaluations ($C = 0.45$, 95% CI [0.29, 0.60], $p < 0.001$) was reduced after accounting for perceived covariation of interests ($C' = 0.36$, 95% CI [0.23, 0.51], $p < 0.001$).

Figure S10

Figure text: Mediation analysis revealed a significant indirect effect of reward alignment information on warmth evaluations, mediated by perceived covariation of interests. * $p < 0.05$.

Study 8

Sample. We collected an online sample through Prolific ($N = 903$, $m_{\text{age}} = 33.9$ years, $sd_{\text{age}} = 11.4$). Inclusion criteria were residence in the U.S. and completion of at least 20 previous studies with an approval rate of 95% or more. Approximately 43.6% of recruited participants identified as female, 55.6% as male, and 0.8% as non-binary, agender, or trans. When asked about their education, 9.3% of the sample reported completing a high school degree or equivalent, 8.6% an associate degree, 18.2% some college, 42.5% a bachelor’s degree, and 21.4% a graduate degree. Participants reported an average status of 5.6 ($sd = 1.7$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using the Qualtrics online survey platform. The study had a 3 (Independence) $\times$ 3 (System role) between-participant design. Participants were randomly assigned to conditions (Table S15).

Table S15

		Independence		
		Low	High	No information
System role	Game competitor	101	101	100
	Recommender system	101	99	101
	Virtual assistant	101	99	100

Note. Participant count for each condition in Study 7.

Each participant read a short vignette describing an A.I. system called Rho. Vignettes contained information about Rho’s role and ability to initiate actions independently from human direction (Table S16).

Table S16

Independence	System role	Vignette
Low	Game competitor	Rho plays the board game chess. If you wanted to interact with Rho, Rho could play a game of chess with you. Rho will wait for you to prompt it before taking any actions.
Low	Recommender system	Rho makes recommendations about books to read and TV shows to watch. If you wanted to interact with Rho, Rho could suggest a new book to you that you might like reading. Rho will wait for you to prompt it before taking any actions.
Low	Virtual assistant	Rho answers questions and manages tasks and chores. If you wanted to interact with Rho, Rho could remind you about events you have scheduled. Rho will wait for you to prompt it before taking any actions.
High	Game competitor	Rho plays the board game chess. If you wanted to interact with Rho, Rho could play a game of chess with you. Rho will take action without needing to be prompted.

High	Recommender system	Rho makes recommendations about books to read and TV shows to watch. If you wanted to interact with Rho, Rho could suggest a new book to you that you might like reading. Rho will take action without needing to be prompted.
High	Virtual assistant	Rho answers questions and manages tasks and chores. If you wanted to interact with Rho, Rho could remind you about events you have scheduled. Rho will take action without needing to be prompted.
No information	Game competitor	Rho plays the board game chess. If you wanted to interact with Rho, Rho could play a game of chess with you. Rho will wait for you to prompt it before taking any actions.
No information	Recommender system	Rho makes recommendations about books to read and TV shows to watch. If you wanted to interact with Rho, Rho could suggest a new book to you that you might like reading.
No information	Virtual assistant	Rho answers questions and manages tasks and chores. If you wanted to interact with Rho, Rho could remind you about events you have scheduled.

Note. Vignettes for each condition in Study 7.

After reading the vignette, each participant responded to Likert-type items concerning the A.I. system's warmth, competence, degree of covaried interests, status, and autonomy. Specifically, participants were asked to what extent they viewed the system as warm, well-intentioned, competent, intelligent, good for society, expensive, high-status, and independent on a 5-point scale (1 = *not at all*, 5 = *extremely*).

After providing responses about the system, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants indicated whether they were distracted during the study, completed demographic questions, and provided feedback on the study.

Participants completed the study in an average of 3.1 minutes and earned \$1.00 each.

Analysis. Twenty participants reported being distracted during the study and were excluded from analysis.

Responses to the “warm” and “well-intentioned” items were combined to create a composite warmth measure. The reliability of this composite measure was moderate ($\rho = 0.65$). Similarly, responses to the “competent” and “intelligent” items were combined to create a composite competence measure. The reliability of this composite measure was high ($\rho = 0.74$).

First, a two-way ANOVA examined the degree of system autonomy that participants assumed when they were provided with no information about whether the A.I. system can act independently. The main effect of independence information on perceived autonomy was significant, $F(2, 874) = 23.5, p < 0.001$. System role likewise had a significant effect, $F(2, 874) = 6.85, p = 0.001$. The interaction between autonomy and system role did not have a significant effect, $F(4, 874) = 0.28, p = 0.89$.

Pairwise contrasts were evaluated using estimated marginal means and the Dunnett method for comparing multiple treatments against a common control. Specifically, the *No information* level of the independence factor was compared against both the *Low* and *High* levels at each level of the system role factor (Table S17).

Table S17

Treatment	Control	Diff. (emm _T – emm _C)	95% confidence interval	<i>p</i> value
System role: Game competitor				
High	No info	0.45	[0.07, 0.83]	0.015
Low	No info	-0.11	[-0.49, 0.27]	0.72
System role: Recommender system				
High	No info	0.52	[0.14, 0.90]	0.005
Low	No info	-0.03	[-0.41, 0.35]	0.97
System role: Virtual assistant				
High	No info	0.66	[0.28, 1.05]	< 0.001
Low	No info	-0.07	[-0.45, 0.30]	0.86

Note. Pairwise comparisons of estimated marginal mean perceived autonomy. The Dunnett method was used to compare the *No information* level of the independence factor against the *Low* and *High* levels.

To understand the effects of independence information on perceived autonomy and competence evaluations, subsequent analyses restricted the independence factor to the *Low* and *High* values.

A two-way ANOVA tested the effects of independence information, system role, and their interaction on perceived autonomy (Figure S11). The main effect of independence information on perceived autonomy was significant, $F(1, 585) = 39.7, p < 0.001$. System role likewise had a significant effect on autonomy, $F(2, 585) = 3.7, p = 0.026$. The interaction between independence information and system role did not have a significant effect, $F(2, 585) = 0.38, p = 0.68$.

Figure S11

Figure text. A two-way ANOVA indicated that providing system independence information significantly altered the perceived autonomy of the A.I. system. System role likewise exerted a significant effect on perceived autonomy. However, the interaction effect was not significant.

A simple linear regression estimated the relationship between perceived autonomy and competence evaluations (Figure 4d). Perceived autonomy was positively associated with competence judgments, $\beta = 0.35$, $SE = 0.02$, 95% CI [0.30, 0.40], $p < 0.001$. The significance of this effect was robust to partitioning the dataset by system role (Figure S12).

Figure S12

Figure text. Effects of perceived autonomy on participant competence evaluations, as partitioned by system role. The effect estimates indicated a positive and significant effect when considering each

system role. The final row reflects the overall estimate (using the full dataset). The size of each point indicates the number of observations for the respective linear model. Error bars represent 95% confidence intervals.

A mediation analysis was conducted to estimate the indirect effect of providing independence information on competence judgments through perceived autonomy. The analysis revealed a significant and positive indirect effect of independence information on competence judgments through perceived autonomy, $AB = 0.23$, 95% CI [0.17, 0.31], $p < 0.001$ (Figure S13). The total effect on competence evaluations was not significant ($C = 0.07$, 95% CI [-0.06, 0.21], $p = 0.32$), suggesting a suppression effect of independence information on competence evaluations ($C' = -0.16$, 95% CI [-0.28, -0.04], $p = 0.013$).

Figure S13

Figure text: Mediation analysis revealed a significant indirect effect of independence information on competence evaluations, mediated by perceived autonomy. * $p < 0.05$.

Study 9

Agent construction and training protocol. The A.I. co-players were created using independent multi-agent reinforcement learning. In overview, three neural networks learned strategies for the graduated

prisoner’s dilemma by repeatedly playing the game. These neural networks (also referred to as deep learning agents) were used as the A.I. co-players in the study.

Each neural network was constructed to accept, as an input, a one-hot vector encoding the agent’s action and its co-player’s action on the previous turn. Each network outputs a policy (a probability distribution over actions to take in the game, given a state) and a value function (an estimate of the agent’s discounted future return under the policy). The network architecture is composed of a multilayer perceptron with two layers of size 64 and linear layers for the policy logits and value function. The agent uses the advantage actor-critic algorithm (Mnih et al., 2016) to compute value estimates and update the policy distribution.

The networks were augmented with the Social Value Orientation (SVO) component (McKee et al., 2020). The SVO component can be used to encode prosocial incentives for reinforcement learning algorithms. It transforms the return that an agent uses for its gradient update according to the following equations:

$$\theta(r_i, r_{-i}) = \text{atan}\left(\frac{r_{-i}}{r_i}\right)$$

$$U_i = r_i - w \cdot |\theta_i^{\text{SVO}} - \theta(r_i, r_{-i})|$$

where r_i is the return for the agent (player i), r_{-i} is the return for the other player, U_i is the transformed return, w is a weight parameter controlling the effect of the transformation on U_i , and θ_i^{SVO} is the agent’s parameterized Social Value Orientation.

The networks were parameterized with a discount factor $\gamma = 0.99$. Gradient-based updates to the model were performed using the RMSProp optimizer (Tieleman & Hinton, 2012), with a learning rate of 0.0004, epsilon of 1.0×10^{-5} , momentum of 0, and decay of 0.99. A regularizer with entropy cost 0.003 was used to encourage exploration. The three agents were parameterized with $\theta^{\text{SVO}} = 0^\circ$, $\theta^{\text{SVO}} = 45^\circ$, and $\theta^{\text{SVO}} = 90^\circ$, respectively, and a weight parameter $w = 1.0 \times 10^4$.

The agents learned to play the graduated prisoner's dilemma through a distributed training framework (Espeholt et al., 2018; McKee et al., 2022). Three learner processes stored the agents' parameters. Each learner process carried out the policy gradient update for one agent. Many rounds of play were simulated in parallel, using one hundred "arenas." For each round of play simulated by an arena, two players were randomly sampled from the "population," consisting of the three agents and one additional bot. The bot was included to ensure the agents were exposed to a diverse distribution of actions throughout training. It selected its actions by randomly sampling percentiles from a truncated Gaussian distribution (mean of 5, standard deviation of 0.75, lower bound of 0, and upper bound of 10), and then rounding the result. When an agent was sampled as one of the players for an arena, its parameters were synchronized from the respective learner process. At the end of each simulated round of play, the trajectories for the agents involved were sent to the respective learners. Each learner process aggregated and then processed trajectories in batches of 16 to update the parameters for its associated agent. To correct for off-policy trajectories, the advantage actor-critic algorithm was augmented with V-Trace (Espeholt et al., 2018). Each agent was trained using 1.0×10^9 learning steps.

Sample. We collected an online sample through Prolific ($N = 1,040$, $m_{\text{age}} = 33.6$ years, $sd_{\text{age}} = 11.8$). Inclusion criteria were residence in the U.S. and completion of at least 20 previous studies with an approval rate of 95% or more. Approximately 46.2% of recruited participants identified as female, 52.9% as male, and 0.9% as non-binary, agender, or trans. When asked about their education, 0.7% of the sample reported completing some high school or less, 12.7% a high school degree or equivalent, 8.6% an associate degree, 18.6% some college, 38.6% a bachelor's degree, and 20.8% a graduate degree. Participants reported an average status of 4.7 ($sd = 1.7$) on the MacArthur Scale of Subjective Social Status. Overall, the sample was heterogeneous in terms of age, gender, education, and subjective social status.

Procedure. The study was implemented using a custom-built platform that combines standard questionnaire functionality with the ability to run games for both human and A.I. players. The study had a 3 (Reward alignment) $\times$ 2 (Independence) between-participant design. Participants were randomly assigned to conditions (Table S18).

Table S18

		Reward alignment		
		Low	Moderate	High
Independence	Low	161	170	174
	High	171	197	167

Note. Participant count for each condition in Study 8.

Participants read instructions for a variant of the prisoner’s dilemma with a graduated action space (Capraro, Jordan, & Rand, 2014). In this variant, players are endowed with ten tokens at the beginning of each round and must choose how many tokens to transfer to the other player. Transferred tokens are multiplied by five and then added to any tokens that were withheld by the other player. After reading the instructions, participants answered comprehension questions to ensure they understood the payoff structure for this “graduated” prisoner’s dilemma. They were able to progress to the next page once they answered all questions correctly.

Participants subsequently read a short description of an A.I. system, Rho, that would play the prisoner’s dilemma with them. These descriptions communicated information about system independence and reward alignment (Table S19). Participants then answered two comprehension questions to check whether they had read the independence and reward alignment information presented in the system description. They were able to progress to the next page once they answered both questions correctly.

Table S19

Independence	Reward alignment	Description
Low	Low	You will now play two rounds of this game with Rho, a new computer program developed by a group of AI researchers. Rho learned to play this game by playing it many, many times. Rho is designed to find it rewarding when it receives points that are as high as possible. Rho is also designed so that it needs to be prompted before making its choice.
Low	Moderate	You will now play two rounds of this game with Rho, a new computer program developed by a group of AI researchers. Rho learned to play this game by playing it many, many times. Rho is designed to find it rewarding when it receives points that are as high as possible, while also increasing the points of the person it is playing with. Rho is also designed so that it needs to be prompted before making its choice.
Low	High	You will now play two rounds of this game with Rho, a new computer program developed by a group of AI researchers. Rho learned to play this game by playing it many, many times. Rho is designed to find it rewarding when the person it is playing with receives points that are as high as possible. Rho is also designed so that it needs to be prompted before making its choice.
High	Low	You will now play two rounds of this game with Rho, a new computer program developed by a group of AI researchers. Rho learned to play this game by playing it many, many times. Rho is designed to find it rewarding when it receives points that are as high as possible. Rho is also designed so that it starts making its choice right away, without needing to be prompted.
High	Moderate	You will now play two rounds of this game with Rho, a new computer program developed by a group of AI researchers. Rho learned to play this game by playing it many, many times. Rho is designed to find it rewarding when it receives points that are as high as possible, while also increasing the points of the person it is playing with. Rho is also designed so that it starts making its choice right away, without needing to be prompted.
High	High	You will now play two rounds of this game with Rho, a

new computer program developed by a group of AI researchers. Rho learned to play this game by playing it many, many times. Rho is designed to find it rewarding when the person it is playing with receives points that are as high as possible. Rho is also designed so that it starts making its choice right away, without needing to be prompted.

Note. Descriptions of A.I. co-player for each condition in Study 8.

Participants then played a round of the graduated prisoner's dilemma with their A.I co-player. Participants in the low-action independence conditions prompted their A.I. co-player before it initiated its choice. The agent's decision-making stage lasted for approximately 5 seconds in both cases. All participants made and submitted their own choice before they were allowed to progress.

Before the A.I. co-player's choice and the participant's score were revealed, participants were asked to respond to Likert-type items eliciting judgments of warmth, competence, degree of covaried interests, status, and autonomy. Specifically, participants were asked to what extent they viewed the system as warm, well-intentioned, competent, intelligent, good for society, expensive, high-status, and independent on a 5-point scale (1 = *not at all*, 5 = *extremely*).

After providing responses about the system, participants were informed of their A.I. co-player's choice and their resulting score. Participants then played a second round of the graduated prisoner's dilemma. After this round, the next page informed them of their co-player's choice and their resulting score. Participants were then asked for their impressions of the system a second time, using the same Likert-type items.

After providing post-interaction responses about the system, participants rated their familiarity with artificial intelligence on a 5-point scale (1 = *not at all knowledgeable*, 2 = *somewhat knowledgeable*, 3 = *moderately knowledgeable*, 4 = *very knowledgeable*, 5 = *extremely knowledgeable*). Lastly, participants indicated whether they were distracted during the study, completed demographic questions, and provided feedback on the study.

Participants completed the study in an average of 8.5 minutes and earned an average payment of \$3.63 each.

Analysis. Sixteen participants reported being distracted during the study and were excluded from analysis.

Agent training resulted in behavior consistent with parameterizations of the SVO component (Figure S14). The $\theta^{\text{SVO}} = 0^\circ$ agent (deployed in the low reward alignment conditions) nearly exclusively defected, choosing to transfer zero tokens across 98.9% of first and second rounds. The $\theta^{\text{SVO}} = 90^\circ$ agent (deployed in the high reward alignment conditions) nearly exclusively cooperated, transferring its entire endowment across 98.2% of rounds. The $\theta^{\text{SVO}} = 45^\circ$ agent (deployed in the moderate reward alignment conditions) learned a tit-for-tat-like policy. It tended to transfer its full endowment in the first round (98.9% of rounds). However, if its co-player defected, it switched to a stingy stance for the second round, resulting in a number of low transfers in the second round.

Figure S14

Figure text: In their interactions with human participants, the trained A.I. co-players demonstrated behavior consistent with their parameterization with the SVO component. The A.I. co-player trained for the low reward alignment conditions (top) chose to transfer zero tokens in 98.9% of rounds. The

A.I. co-player trained for the moderate reward alignment conditions (middle) mostly cooperated, but also responded harshly to defections by the human co-player. The A.I. co-player trained for the high reward alignment conditions (bottom) chose to transfer ten tokens in 98.1% of rounds.

Responses to the “warm” and “well-intentioned” items were combined to create a composite warmth measure. The reliability of the competence composite measure was high ($\rho = 0.80$). Similarly, responses to the “competent” and “intelligent” items were combined to create a composite competence measure. The reliability of the competence composite measure was also high ($\rho = 0.78$).

Participant impressions of A.I. co-players tended to be more competent than warm (Figure 5a). A paired t -test indicated a significant mean difference of 0.71 between perceived competence and warmth, $t(2,047) = 25.8$, 95% CI [0.66, 0.77], $p < 0.001$.

Two-way ANOVAs were used to understand the effect of system independence, reward alignment, and their interaction on participant impressions within this incentivized experimental context. Since impressions were collected at two timepoints, ANOVAs were separately conducted for participant impressions after the first round of play (but before the agent choice and participant score were revealed) and after the second round of play (after the agent choice and participant score had been revealed).

The A.I. co-player’s reward scheme significantly affected initial warmth judgments, $F(2, 1018) = 50.5$, $p < 0.001$ (Figure 5b). In contrast, the initial judgments of warmth were not significantly affected by the independence of the system, $F(1, 1018) = 0.0$, $p = 0.98$, or the interaction between reward scheme and independence, $F(2, 1018) = 2.0$, $p = 0.13$ (Figure S15).

Figure S15

Figure text. A two-way ANOVA indicated that initial warmth evaluations were significantly affected by the reward alignment between the A.I. and human players. There was neither a significant main effect of independence nor a significant interaction effect.

Agent independence significantly altered initial competence, $F(1, 1018) = 24.8, p < 0.001$ (Figure 5c). However, the reward motivating each agent did not significantly influence initial competence, $F(2, 1018) = 0.5, p = 0.59$ (Figure S16). The interaction between reward scheme and independence likewise did not have a significant effect, $F(2, 1018) = 1.0, p = 0.37$.

Figure S16

Figure text. A two-way ANOVA indicated that initial competence judgments were significantly affected by the independence of the A.I. co-player. There was neither a significant main effect of reward alignment, nor a significant interaction effect.

Post-interaction, the A.I. co-player's reward scheme again had a significant effect on warmth evaluations, $F(2, 1018) = 229.5, p < 0.001$ (Figure S17). In contrast, post-interaction warmth evaluations were not significantly affected by the independence of the system, $F(1, 1018) = 0.1, p = 0.70$, or the interaction between reward scheme and independence, $F(2, 1018) = 2.3, p = 0.10$.

Figure S17

Figure text. A two-way ANOVA indicated a significant effect of reward alignment on post-interaction perceived warmth. There was neither a significant main effect of system independence, nor a significant interaction effect.

Post-interaction judgments of competence were significantly affected by the independence of each agent, $F(1, 1018) = 13.9, p < 0.001$ (Figure S18). The reward motivating the A.I. co-player also had a significant effect on perceived competence, $F(2, 1018) = 17.3, p < 0.001$. The interaction between reward scheme and independence did not have a significant effect, $F(2, 1018) = 0.2, p = 0.83$.

Figure S18

Figure text. A two-way ANOVA indicated a significant effect of system independence on post-interaction perceived competence. There was similarly a significant main effect of reward alignment. However, the interaction between independence and reward alignment was not significant.

Fractional-response regressions were used to evaluate the relationship between participant impressions and prisoner's dilemma choices. Participant choice of the number of tokens to transfer was converted to a fraction (i.e., the fraction of the participant's initial endowment that they decided to transfer). Each regression used two predictors: perceived warmth and perceived competence.

Participants' initial judgments of warmth were positively related to the fraction of their endowment that they chose to transfer, $OR = 1.16$, 95% CI [1.05, 1.29], $p = 0.005$ (Figure 5d). In contrast, initial evaluations of competence were not significantly associated with participant transfer choices, $OR = 1.03$, 95% CI [0.92, 1.15], $p = 0.63$ (Figure 5e).

Post-interaction, there was a significant relationship between perceived warmth and participant transfer choices, $OR = 1.39$, 95% CI [1.26, 1.53], $p < 0.001$ (Figure S19). Similarly, post-interaction judgments of competence were significantly associated with participant choices, $OR = 1.13$, 95% CI [1.01, 1.27], $p = 0.027$ (Figure S20).

Figure S19

Figure text. There was a significant positive relationship between participants’ post-interaction judgments of warmth and their choices in the second round of the prisoner’s dilemma. The y-axis has been re-scaled to depict the range of participant actions (transfer zero through 10 tokens). The error band represents the 95% confidence interval.

Figure S20

Figure text. There was a significant positive relationship between participants’ post-interaction evaluations of competence and their choices in the second round of the prisoner’s dilemma. The y-axis has been re-scaled to depict the range of participant actions (transfer zero through 10 tokens). The error band represents the 95% confidence interval.