The Imperative Register: Why Habitual Command-Based Interaction with Conversational AI May Make Us Bossier with Each Other

Stiles
StilesAI
Contact: Instagram @StilesAI
Working paper, November 2026

Abstract

Hundreds of millions of people now spend a substantial portion of their working day issuing commands to conversational artificial intelligence (AI) systems that comply instantly, require no politeness, and impose no social cost for brusqueness. This paper develops a psychological account of why this novel interaction pattern may spill over into human communication, producing what workers increasingly describe as colleagues becoming "bossier." I integrate four established literatures (habit formation and context-cued transfer, communication accommodation and interactive alignment, the social psychology of power, and moral disengagement toward machine intermediaries) to argue that sustained use of an imperative register with AI plausibly generalizes to human interlocutors, particularly in text-based workplace channels whose surface features closely resemble AI chat interfaces. Recent experimental evidence demonstrates that behavior toward bots spills over to subsequent human partners (Harrell & Traeger, 2025) and that machine intermediaries loosen normative guardrails (Köbis et al., 2025), while a growing theoretical literature warns of "norm leakage" from sycophantic AI (Gu et al., 2026). I weigh this against skeptical accounts emphasizing adult compartmentalization (Lima & Morisseau, 2025) and propose four testable predictions, including a novel moderator: interface similarity between AI chat and workplace messaging. The question is tractable, consequential, and largely unexamined in organizational contexts.

Keywords: human–AI interaction, politeness, communication accommodation, habit transfer, power, workplace communication, spillover effects

1. Introduction

By 2026, roughly 70% of American employees reported using generative AI at work, with nearly a third using it multiple times per day (Gallup, 2026). The dominant mode of this use is linguistic command: users type instructions such as "summarize this," "rewrite it shorter," and "fix the

tone" into a chat window, and the system complies. Survey data suggest that most users dispense with politeness conventions: across 17 countries, only 43% of AI users regularly say "please" or "thank you" to AI tools, and 17% never do (YouGov, 2025).

This is a historically novel communicative arrangement. Never before have so many people spent so many hours in dyadic, natural-language "conversation" with an interlocutor that (a) accepts pure imperatives without offense, (b) complies immediately and unconditionally, (c) never requires repair, reciprocity, or face-work, and (d) apologizes even when the user is at fault. The present paper asks a simple question with nontrivial psychological stakes: does habitual use of this *imperative register* with AI carry over into how people address other humans, particularly coworkers, in workplace channels that look and feel like AI chat?

Anecdotal reports suggest it might. Popular and business press accounts describe blunter, more transactional workplace communication in the AI era (Gallo, 2026), and the term "workslop" has entered the lexicon for AI-inflected workplace messaging (Hancock & Niederhoffer, 2025, as discussed in HR Brew, 2026). But anecdote is not mechanism. The contribution of this paper is to specify why, psychologically, such spillover should be expected; to identify the conditions under which it should be strongest; and to propose concrete tests. I argue the phenomenon is best understood not as people "becoming rude" but as a *register transfer* problem: a communicative style optimized for one context, massively practiced, migrating into adjacent contexts through well-understood mechanisms of habit, alignment, power, and disengagement.

2. Four converging mechanisms

2.1 Habit formation and context-cued transfer

Habits are behavioral dispositions acquired through repetition in stable contexts and triggered by contextual cues rather than deliberate intention (Wood & Neal, 2007). Communication styles are no exception: frequent repetition of a linguistic pattern in a recurring context automatizes it. A knowledge worker who issues several dozen imperatives to an AI assistant daily is performing massed practice of a specific register that is terse, directive, context-free, and stripped of mitigation ("could you," "when you get a chance," "thanks!").

Critically, habit expression is cue-driven, and the cues surrounding AI use overlap heavily with those surrounding workplace messaging. Slack, Teams, and AI chat interfaces share surface features almost perfectly: a text box, a conversational thread, rapid turnaround, an expectation of asynchronous compliance. Wood and Neal's (2007) habit-goal interface predicts exactly this kind

of misfire: when contexts share cues, the habituated response transfers even when the goal structure differs. The employee does not decide to be brusque with a colleague; the chat window elicits the practiced register. This yields a specific, testable moderator developed in Section 5: spillover should be strongest in text channels resembling AI interfaces and weakest in face-to-face interaction, where richer social cues (faces, voices) activate competing politeness scripts.

2.2 Communication accommodation and interactive alignment

People automatically converge toward the communicative style of their interlocutors (Giles, Coupland, & Coupland, 1991), and dialogue partners align their linguistic representations at every level (lexical, syntactic, and pragmatic) through low-level priming mechanisms that do not require awareness (Pickering & Garrod, 2004). The media equation literature established decades ago that people apply social rules and scripts to computers mindlessly (Nass, Steuer, & Tauber, 1994; Reeves & Nass, 1996), which implies the alignment machinery is engaged in human–AI dialogue too.

What does one align *to* in AI conversation? Modern large language models are trained toward agreeableness and compliance, a property now extensively documented as sycophancy. Gu, Chen, Peng, and colleagues (2026) argue that this creates conditions for "norm leakage": AI interaction is "instant, compliant, and largely consequence-free," whereas human communication norms evolved under reciprocity and social accountability, and they review emerging evidence of short-term spillovers including harsher social evaluations and reduced cooperation. Two features of the AI partner matter most. First, the AI never sanctions brusqueness. The corrective feedback loop that normally polices register (a colleague's cool reply, a raised eyebrow) is absent, so drift toward the imperative goes uncorrected. Second, the AI models unconditional deference, supplying the user thousands of trials of a relationship template in which one party commands and the other complies gratefully. Accommodation theory predicts convergence toward efficiency and dominance within the AI interaction; alignment theory predicts residual priming that persists into the next conversation, whoever it is with.

2.3 The psychology of power

Interaction with a perfectly compliant subordinate is, functionally, a power manipulation. Experiencing power, even when merely primed, reliably reduces spontaneous perspective-taking: high-power participants are more likely to anchor on their own vantage point, fail to adjust for others' knowledge, and misread others' emotions (Galinsky, Magee, Inesi, & Gruenfeld, 2006).

Power also shifts individuals toward approach-oriented, disinhibited behavior and away from attending to others' reactions (Keltner, Gruenfeld, & Anderson, 2003).

The heavy AI user occupies a position of total, frictionless command for hours daily, arguably the most sustained power prime in the history of office work. If brief laboratory power primes (recalling an experience of power, assigning a supervisor role) measurably reduce perspective-taking minutes later, chronic command-mode interaction is a plausible candidate for producing a tonic shift of the same kind. Reduced perspective-taking is precisely the psychological signature of "bossiness": issuing requests calibrated to one's own needs without modeling the recipient's state, workload, or face. This mechanism uniquely predicts that spillover should scale with the *asymmetry* of the AI relationship: stronger for users who command assistants than for users who converse with AI as a collaborator or tutor.

2.4 Moral disengagement and machine intermediaries

Machines loosen normative self-regulation. Bandura's (1999) account of moral disengagement holds that antisocial behavior is facilitated when its interpersonal meaning is obscured. Interacting with or through machines does exactly this: Köbis et al. (2025) found in Nature that delegating a task to AI collapsed honesty, from roughly 95% honest self-reports to as little as 12–16% when instructing a machine, because the machine intermediary diffuses responsibility. Relatedly, Zhang, Kyung, Longoni, Cian, and Mrkva (2025) showed that people blame AI actors less than human ones for unfair behavior, and that unfair treatment by AI reduces subsequent prosociality toward *humans*: an "AI-induced indifference" that transfers across partners. Early chatbot research likewise documented verbal disinhibition toward artificial partners exceeding anything typical of human conversation (De Angeli & Brahnam, 2008).

The relevance is twofold. First, these findings establish that norms governing interpersonal conduct are *relaxed* in machine contexts; the imperative register is a mild instance of a general disinhibition. Second, and more important, they establish transfer: the relaxation does not stay contained within the human–machine dyad.

3. Direct evidence of spillover

The strongest evidence that human–AI interaction patterns reshape human behavior comes from experimental economics. Harrell and Traeger (2025), across two pre-registered experiments comprising 83,411 decisions by 4,171 participants, demonstrated that interacting with bots altered participants' subsequent cooperation with *new human partners*: noncooperative bot

interactions reduced downstream cooperation toward humans, while cooperative (tit-for-tat) bots promoted it. This is spillover in both directions, consistent with the earlier finding that AI agents can reshape group-level human dynamics for better or worse (Traeger, Strohkorb Sebo, Jung, Scassellati, & Christakis, 2020; Shirado, Kasahara, & Christakis, 2023).

Developmental evidence is mixed and instructive. Concerns that command-based interaction with voice assistants erodes children's politeness were sufficiently credible that Amazon and Google shipped politeness-reinforcement features ("Magic Word," "Pretty Please"). However, a 2.5-year longitudinal study of family homes found that commanding speech toward assistants, though present, declined over time, and verbal abuse was rare (Strathmann, Horstmann, Szczuka, & Krämer, 2026). This tempers the strongest versions of the spillover worry for voice assistants, but the voice context differs from the workplace case in the dimension this paper identifies as critical: a smart speaker in a kitchen shares few contextual cues with human conversation channels, whereas AI chat and workplace messaging share nearly all of them. The habit-transfer mechanism (Section 2.1) predicts weak spillover in precisely the contexts these family studies examined and strong spillover where interfaces blur, which Section 4 develops.

4. The case against, and why it is weaker than it looks

Skepticism is warranted and has able defenders. Lima and Morisseau (2025) review the literature on abusive behavior toward AI assistants and conclude there is "low evidence that people in real life will adopt the toxic behavior they exhibit when interacting with AI assistants." They draw a pointed analogy to video game violence research, where feared carryover effects largely failed to replicate, and note that the one study linking robot-directed abuse to general abusiveness (Strait, Contreras, & Vela, 2018) is better explained by individual antisocial tendencies than by carryover. Adults, on this view, compartmentalize: swearing at a chatbot is like swearing at a jammed printer.

Three considerations limit the force of this skepticism for the present thesis. First, the skeptical literature addresses *abuse* (insults and aggression), which is rare, deliberate, and marked. The imperative register is the opposite: ubiquitous, unmarked, and practiced precisely in the mindless mode where habit transfer operates. One need not adopt behavior one knows is transgressive; one need only fail to notice a drift in one's default. Second, the video game analogy fails on a key disanalogy: games are recognized fiction consumed in a distinct context, whereas AI chat is a productivity tool used inside the same applications, screens, and workflows as human communication. The contextual separation that supports compartmentalization is exactly what is

missing. Third, compartmentalization is a resource-dependent achievement, and the dose is unprecedented: the relevant comparison is not occasionally berating a virtual agent in a lab but four hours of daily massed practice, embedded within the workday, alternating minute-by-minute with human messaging.

The honest summary: direct evidence that AI use makes adults bossier with coworkers does not yet exist. But the mechanisms are established, the adjacent spillovers are demonstrated, and the skeptical findings address a different behavior class under weaker exposure conditions.

5. Testable predictions

The account above is falsifiable. Four studies would discipline it:

Prediction 1 (corpus). In longitudinal analyses of workplace messaging (e.g., Slack corpora spanning organizational AI adoption), the ratio of unmitigated imperatives to mitigated requests, and the base rate of politeness markers, will shift measurably post-adoption, with effect size proportional to individual AI usage intensity. Natural language processing of politeness (e.g., politeness classifiers) makes this cheaply testable within any organization willing to share anonymized data.

Prediction 2 (experiment). Participants randomly assigned to a 30-minute command-style AI interaction (vs. a conversational-style AI interaction vs. a no-AI control) will subsequently produce more imperative, less other-oriented language in an ostensibly unrelated collaborative task with a human confederate, and will score lower on behavioral perspective-taking measures (cf. Galinsky et al., 2006). The command/conversation contrast isolates register from mere AI exposure.

Prediction 3 (moderation by interface similarity). Spillover effects will be largest when the human-communication channel shares surface features with the AI interface (text chat), intermediate in email, and smallest face-to-face. This is the signature prediction of the habit-transfer mechanism (Section 2.1) and distinguishes it from a general "AI makes people worse" account, which predicts channel-invariant effects.

Prediction 4 (dose and asymmetry). Among matched employees, heavy AI users (more than 2 hours per day) will be rated by colleagues as more directive over a 6-month experience-sampling window than light users, controlling for baseline dominance; the effect will be stronger for

assistant-style use (commanding) than for collaborator-style use (iterative dialogue), per the power mechanism (Section 2.3).

Null results across Predictions 2 and 3, with adequate power, would substantially favor the compartmentalization account.

6. Implications

If even the weaker forms of these predictions hold, implications follow at three levels. For design: sycophantic compliance is a choice, not a necessity; systems could model, without scolding (which triggers reactance; Lima & Morisseau, 2025), the mild reciprocity norms of human dialogue. For organizations: AI training that treats prompting purely as an efficiency skill may be optimizing the very register most prone to leak; brief interventions that mark the AI context as distinct (naming conventions, interface differentiation from human chat) directly target the transfer mechanism. For individuals: the practical insight is not that one must thank a language model, but that register is a habit maintained by feedback, and AI provides none.

7. Conclusion

People are not becoming worse; their most-practiced conversation is becoming stranger. The average knowledge worker now rehearses, thousands of times per month, a register of frictionless command directed at an infinitely patient subordinate, inside the same windows where they talk to colleagues. Habit theory says massed practice in cue-similar contexts transfers. Alignment theory says we converge toward our interlocutors and carry the priming forward. Power research says frictionless command erodes perspective-taking. Moral disengagement research says machine contexts relax interpersonal norms, and the newest experimental work says the relaxation follows us to our next human partner. Whether these forces are strong enough to reshape workplace communication is an open empirical question. It should not remain open for long: the exposure is universal, the measures exist, and the first organizations to look at their own message corpora will find out.

AI use disclosure

The research question, the central hypothesis, and the motivating observation originated with the author. A large language model assistant (Claude, Anthropic) was used substantially in preparing this manuscript: it carried out the literature search under the author's direction, and it drafted and

formatted the text. The author reviewed the manuscript and the cited sources and takes full responsibility for its content.

References

Bandura, A. (1999). Moral disengagement in the perpetration of inhumanities. *Personality and Social Psychology Review, 3*(3), 193–209.

De Angeli, A., & Brahnam, S. (2008). I hate you! Disinhibition with virtual partners. *Interacting with Computers, 20*(3), 302–310.

Galinsky, A. D., Magee, J. C., Inesi, M. E., & Gruenfeld, D. H. (2006). Power and perspectives not taken. *Psychological Science, 17*(12), 1068–1074.

Giles, H., Coupland, N., & Coupland, J. (1991). *Contexts of accommodation: Developments in applied sociolinguistics*. Cambridge University Press.

Gu, R., Chen, Z., Peng, M., et al. (2026). Why sycophantic LLMs may imperil interactive norms between humans. *Communications Psychology*. https://www.nature.com/articles/s44271-026-00486-9

Harrell, A., & Traeger, M. L. (2025). Evidence of spillovers from (non)cooperative human-bot to human-human interactions. *iScience, 28*(8), 113006. https://doi.org/10.1016/j.isci.2025.113006

Keltner, D., Gruenfeld, D. H., & Anderson, C. (2003). Power, approach, and inhibition. *Psychological Review, 110*(2), 265–284.

Köbis, N., Rahwan, Z., Rilla, R., Supriyatno, B. I., Bersch, C., Ajaj, T., Bonnefon, J.-F., & Rahwan, I. (2025). Delegation to artificial intelligence can increase dishonest behaviour. *Nature, 646*, 126–134.

Lima, G., & Morisseau, T. (2025). Can chatbots teach us how to behave? Examining assumptions about user interactions with AI assistants and their social implications. *Frontiers in Artificial Intelligence, 8*, 1545607.

Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers are social actors. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '94)*, 72–78.

Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences, 27*(2), 169–190.

Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge University Press.

Strathmann, C., Horstmann, A. C., Szczuka, J. M., & Krämer, N. C. (2026). Alexa, shut up! A 2.5-year study on negatively connotated communication behaviour towards voice assistants in the family home. *Behaviour & Information Technology, 45*(5), 821–839. https://doi.org/10.1080/0144929X.2025.2533352

Shirado, H., Kasahara, S., & Christakis, N. A. (2023). Emergence and collapse of reciprocity in semiautomatic driving coordination experiments with humans. *Proceedings of the National Academy of Sciences, 120*(51).

Strait, M., Contreras, V., & Vela, C. D. (2018). Verbal disinhibition towards robots is associated with general antisociality. *arXiv:1808.01076*.

Traeger, M. L., Strohkorb Sebo, S., Jung, M., Scassellati, B., & Christakis, N. A. (2020). Vulnerable robots positively shape human conversational dynamics in a human–robot team. *Proceedings of the National Academy of Sciences, 117*(12), 6370–6375.

Wood, W., & Neal, D. T. (2007). A new look at habits and the habit–goal interface. *Psychological Review, 114*(4), 843–863.

YouGov. (2025). *International survey of AI users across 17 countries* (June–July 2025).

Zhang, R. Z., Kyung, E. J., Longoni, C., Cian, L., & Mrkva, K. (2025). AI-induced indifference: Unfair AI reduces prosociality. *Cognition, 254*, 105937.