

Video Call Glitches Trigger Uncanniness and Harm Consequential Life Outcomes

Corresponding Author: Dr Melanie Brucks

Any redactions in this file are there to maintain patient confidentiality, the confidentiality of unpublished data, or to remove third-party material.

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Attachments originally included by the reviewers as part of their assessment can be found at the end of this file.

Note from the editor: Reviewer 2 was unable to re-review the revised manuscript. However, as some of the other referees raised overlapping concerns about the effect sizes, we asked for their advice on the authors' responses of their revised manuscript and they deemed them satisfactory.

Version 1:

Reviewer comments:

Referee #1

(Remarks to the Author)

The paper presents extensive empirical evidence, both causal and correlational, demonstrating negative effects of glitches in virtual communication on important interpersonal outcomes like liking and trust.

This is a very strong paper. I am impressed by the body of evidence and the rigor of inquiry. Crucially, I think that you have identified an interesting, novel, and important phenomenon. I am especially concerned about the effect of glitches in court hearings. That finding alone merits publication. Below I offer a few comments in the hope that they can help you further strengthen an already strong paper.

1) While the effect documented in the paper has clear relevance for individuals who carry out important meetings remotely (e.g., job interviews), the paper makes a societal claim of substantive importance: the effect of glitches risks entrenching inequalities in society. I am positively predisposed towards this claim, and this is a comment aimed at sharpening this aspect of the paper. A devil's advocate could argue that by and large people across the SES ladder rely on the same data networks nowadays. Is it really true that glitches are more frequently experienced by low SES citizens? For example, to the extent that data connectivity is better in urban centers and that the poor tend to concentrate in urban areas (vs. less densely populated areas surrounding cities), one might even make the opposite prediction. I am not writing this because I believe your contention to be false but only to illustrate the need for validating a key assumption you rely on to claim that glitches exacerbate inequality.

2) I think that the paper is currently a bit hand-wavy on the psychological mechanism. You rely on the construct of "uncanniness" but I struggle to fully understand the nature and outlines of this process. You define it as the reaction accompanying a violation of expectations. Could one argue that this construct is too close to the independent variable? (More akin to a manipulation check than a psychological process?) I think the idea makes intuitive sense but perhaps the paper could conceptualize this construct with greater precision. I'll note that despite the large amount of data, the paper contains relatively little process evidence by moderation (in this sense I found Study 5 very important).

3) Related to the previous, and as you also note in the paper, it is relatively easy to argue for alternative psychological processes. For example, lower rates of parole in virtual court appearances with glitches may be the result of the judge's frustration. Glitches in video calls are frustrating indeed.

4) You treat glitches as "one thing" but there are many types of glitches, and some may be more strongly associated to feelings of uncanniness and to negative interpersonal outcomes. For example, I find lack of synchrony between sound and

image very uncanny. Even if other types of glitches may be objectively involving a much larger information loss, if your hypothesis is correct they may not produce effects that are as negative. Going back to the point above about moderation, one option could be to build a typology of glitches, pretest them for uncanniness, and then conduct a process test via moderation. I like the classification provided for Study 1 (p. 26) and perhaps more could be done in this direction. I am mentioning this as an opportunity to deepen your insights and implications. I do not think that the paper needs necessarily more studies given the amount of data that you already gathered.

5) A smaller comment: could the effect reverse as a function of the topic of conversation? It is a comedy cliché for someone on a call to pretend that the connection is poor when facing scrutiny or a difficult conversation. Could it be that glitches that hurt the ability to communicate may sometimes give “cover” to people who find themselves in a difficult spot?

Congratulations for writing a very impressive paper.

Referee #2

(Remarks to the Author)

This is a robust corpus of studies, they have good power, applied across a variety of contexts that show a large spectrum of real world applications.

My overarching concern stems from how practically significant these effects are. Let's say the mean difference in outcome with glitches is 4.25 vs 4.5 and the effect size is significant. But is this a practically meaningful difference. In other words, is 4.25 actually low enough to be concerned about glitches. Yes glitches, do affect the intended outcome somewhat adversely, but it appears the intended outcome is still being achieved.

The study maybe statistically robust and exhaustive, but it fails to drive the importance of the questions and the context. And glitches ultimately will reduce as technologies and access both improve with reduced costs.

Referee #3

(Remarks to the Author)

This paper examines how audiovisual glitches during video calls impact interpersonal judgments and decisions across a broad range of contexts. The authors demonstrate that seemingly trivial technical disruptions (“glitches”) can significantly impact outcomes in settings such as hiring decisions, medical consultations, and even parole hearings. The main mechanism seems that glitches violate people’s social expectations and elicits a sense of “uncanniness”. Given the novelty of this phenomenon, the societal significance, the robust findings with replications across various settings and measurements approaches, I see this paper as a strong candidate for publication in Nature. I first focus on the main themes to assess suitability for publication at Nature, followed by areas that need some (minor) work to improve the robustness of the authors’ claims (section 4).

1) Significance & Originality of Main Finding:

- The phenomenon is pervasive, and the demonstration of the main effect across multiple domains (telemedicine, hiring decisions, judicial proceedings) offers a compelling fit for Nature.
- The documented prevalence of such glitches (affecting up to 30% of video calls) establishes clearly the paper’s societal significance, with substantial implications for equity and access in virtual communication. These findings have the potential to rethink a range of existing research in psychology, virtual communication, technology access, and potentially the future of public policy.
- I particularly value the significance of the paper with respect to the societal implications, as the research reveals how technical disruptions may systematically disadvantage already marginalized populations, potentially amplifying the digital divide across many nations worldwide.

2) Validity of Findings & Methodology:

- The procedures are well described and all materials are exceptionally well documented.
- The paper elegantly and effectively combines field evidence (S1, S6) with controlled experiments. Overall, the empirical package is strong and the documentation of study materials, data, and analysis code is well done.
- Sample size and statistical power across the entire empirical package is appropriate and especially two studies nicely round out the ecological validity of the authors’ baseline finding (field data of S1 and S6).

3) Conclusions:

- The conclusions are generally well-supported by the data and the authors also rule out various alternative explanations (e.g., mere annoyance or frustration, comprehension issues, etc.) either through existing data, statistical controls, or through the supplementary studies (e.g., setting expectations experiment; joke intervention).

4) Comments & Suggested Improvements:

4.1) More Precise Theory Foundation: The empirical package is solid and the finding is interesting and important. What the authors could carve out more assertively is the precise causal chain and theoretical foundation to link (1) the technical properties of glitches to (2) a distinct psychological process triggered by these disruptions (why uncanniness and not other

mechanisms?), and (3) the specific downstream consequence on interpersonal judgments and decisions. It would be helpful to formulate somewhat sharper conceptual boundaries to differentiate why glitches specifically trigger uncanniness rather than more immediate responses like frustration or competence attributions. I also felt that the authors could situate the theory-foundation more in classic theories of social presence or communication synchrony. This would help explain more precisely the psychological pathway from technical disruption to social judgment and would in turn also help to carve out the specific contribution in light of classic theories of communication synchrony or social presence.

4.2) Deeper Insight on Variation in Glitches & Additional Model Specifications: Looking at Table 2 raises the question whether the effect is truly uniform or if the authors could better exploit variation in glitches. Specifically, the effect sizes look much stronger for video compared to mere audio distortions, while more fleeting vs. more sustained disruptions tend to reduce effect sizes. While this variation could justify the focus on video-based interruptions in the experimental work, it also presents opportunities for a deeper empirical investigation. I think some additional analyses of Studies 1 and 6 would be particularly useful. In Study 1, the authors' didn't model interaction effects between different types of glitches (e.g., combined audio-visual glitches) or temporal sequences that are more vs. less disruptive. For example, I would intuitively expect that it makes a difference whether I experience many short disruptions or a few but long and sustained glitches. In Study 6, the authors could also use a LASSO with the glitch dictionary features as predictors and test specifically which types or sequences of glitches evoke the most negative impact on parole outcomes. Such an analysis has the potential to reveal whether certain types or combinations of glitches are particularly damaging in such high-stakes decisions. To be clear: I consider these additional analyses not absolutely critical for the paper's core contribution, but they would deepen our understanding of potentially critical "glitch patterns" of when and how technical disruptions are most problematic.

4.3) Robustness of Glitch Manipulation: All experimental studies induce frequent and short video distortions (freezing the video while showing the person). This raises the question whether your experimental studies are sensitive or relatively robust to alternative glitch interventions? For example, I kept wondering whether you would produce the same effect using (1) a blackscreen or (2) a pixelated image, instead of the video freeze. If these alternative disruptions produce weaker effect, this would even strengthen the authors' predicted theory and at the same time offer practical guidance on how to deal with it (e.g., immediately block the screen to avoid showing the target person).

4.4) Focus & Positioning: As highlighted earlier, the manuscript is generally well-written. But the opening section would benefit from a more streamlined and precise articulation of the paper's core contribution. The authors should more directly establish that seemingly trivial technical glitches can significantly impact social judgments and consequential decisions across various contexts (hiring, healthcare, parole decisions, etc.). Rather than gradually building to this key main finding, I felt that the introduction could more immediately highlight this paradox and offer a more focused opening. This would not only boost reader engagement but would also carve out the paper's theoretical and practical significance much faster and more direct. The somewhat gradual and piecemeal development of ideas almost dilutes the broad importance of the paper.

Minor:

4.5) Code documentation & Data: Please revise your markdown files to (1) use relative instead of absolute file paths (to make it easier for people to replicate your results) and (2) make sure to use groundhog or other version control packages to make sure your code runs with future R and package updates (write code with a focus on better reproducibility). Finally, please also upload the data and code for S6.

Overall, this is a potentially high-impact paper that is topical, produces a set of theoretically interesting and practically important findings, and has the potential to challenge and rethink existing theories of virtual communication across a wide range of disciplines. The above comments are to help the authors further refine the positioning of the paper and to address a few (mostly minor) points of the current empirical package. Good luck with this interesting and important line of work.

Referee #4

(Remarks to the Author)

This was an interesting paper with some intriguing results. But ultimately I don't think it hits the very high bar required for Nature. So unfortunately I recommended rejecting the paper.

I found four issues:

1) The results are not surprising. Anyone that has used Zoom, Teams, Webex or any other video conferencing technology would know that lower quality video calls are frustrating, and some of that frustration can potentially be transferred to the other partner in the call. The RCTs here involved considerable "glitches" – four or five pauses of up to 2 seconds in a 75 second call, which is a lot. This would involve about 10% of the video call time frozen. As such the results seems somewhat unsurprising in that a repeatedly freezing call renders lower post-call feedback scores. Indeed, I might suspect that, for example, if somebody presented on a Zoom call with very bad lighting, a background with many distractions (like kids fighting), or fuzzy bad audio could have similar impacts.

2) The policy implications are limited. Yes, it is true that individuals with lower incomes likely have worse internet connections on average (although this correlation may not be high, as lower income individuals often live in high-density urban areas which have good connections). As such they may suffer somewhat more here. But the scale of this seems limited. By 2024 most low income individuals are working and living fully in person, with little zoom activity. The people that tend to have remote or hybrid jobs are higher income professionals, who usually have reasonable internet connection (as

this is required for their jobs). So, yes, in the paper they had a nice example from a higher stakes law example, but this was during the pandemic. These kinds of high stakes interactions – like law courts – now do not happen on Zoom. So, I also suspect the “so what” here is limited.

3) Correlation vs Causation. I was concerned with omitted variable bias for studies 1 and 5. The types of people that might report “glitches” or say “freeze” in zoom calls may be more liable to complain in general. Relaxed positive individuals likely complain less. So, their calls are not going to get tagged as “glitchy”. More critical negative individuals probably complain more about glitches so get tagged. But these same negative critical people likely give more critical call evaluations. Hence, it is not surprising that people that complain a lot about call quality also complain a lot about call content, generating a correlation between “glitches” and bad call evaluations. Studies 2 to 4 are more convincing as RCTs, although the magnitude of the induced glitches seem very large for such a short call so it’s more than brief glitches being examined here.

4) Inconsistencies: One final thing is the text and the methods section are at times inconsistent. For example, Study 2 was described in the text as having “four brief (<2 second)” freezes, but in the methods as “five freezes...from 0.5 to 2 seconds in duration”. This seems like a material difference given this is a 75 second video. I would like to know the precise length of all the freezes and the total length of the call, ideally we examples to show this. Also, some of the studies lacked information on how long the freezes lasted for (study 4 for example).

So overall I don’t think this is a good fit for Nature. The Brucks and Levav paper was great – really important results on the impact of video calls on innovation. This paper is less surprising, less well executed, and I think has far narrower policy implications. As such, really sorry, but I don’t think it hits the high bar for Nature.

Version 2:

Reviewer comments:

Referee #1

(Remarks to the Author)

Thank you for being so responsive to my comments! I read the revised paper and the accompanying revision notes with great interest.

In particular, I very much appreciated the new multi-wave study which leverages differences in the uncanniness of various glitch types to demonstrate the effect of glitches more robustly, provide compelling evidence for the role of uncanniness, and suggest important practical implications about the glitch types that require most attention. This is exactly the kind of study I had in mind—in fact, it’s better than that.

You have effectively addressed all my comments. The revised theorizing pinpoints more precisely what feelings of uncanniness are and why they emerge when video calls glitch. Additionally, the new discriminant validity tests reassure about the meaningfulness of the mediation analyses. I have no further comments.

Congratulations for writing a compelling and important paper. It will hopefully help improve decision making in many consequential situations and reduce entrenched inequalities.

Referee #3

(Remarks to the Author)

The authors were very responsive in revising the previous version of the manuscript. They addressed all of my previous comments related to a stronger theoretical foundation rooted in prior work on social presence and media synchrony and also my various methodological comments (e.g., better leveraging and modeling the variation of glitch types and how different types of glitches might produce variance in effect sizes). Although the authors didn’t follow my suggestions on using a LASSO or other regularization techniques, new Study 4 using a multi-wave study provides an excellent and robust improvement with respect to the generalizability of the effect across a variation of glitches. The revised manuscript makes a timely and important contribution to our understanding of technology-mediated communication. While I have a few minor suggestions to help further improve the manuscript, the authors have successfully addressed all of my previous major concerns.

Summary of Key Results

The paper demonstrates that minor audiovisual glitches during video calls systematically harm interpersonal judgments across multiple contexts (hiring, healthcare, parole decisions). The proposed mechanism involves glitches breaking the “illusion of face-to-face contact” that video calls provide, triggering feelings of uncanniness that contaminate subsequent evaluations. The empirical package and evidence for the effect is extensive with ten studies spanning archival datasets and controlled experiments.

Originality and Significance

The core finding—that minor glitches undermine social judgments—is novel and theoretically interesting. Given the ubiquity of video calling post-pandemic, the findings have broad implications for fairness in digital interactions across healthcare, employment, legal proceedings, and social connection more broadly.

Data & Methodology

Overall large sample sizes and one new multi-wave study, using a combination of both archival and experimental data

across the empirical package. Pre-registration of key studies. The data and methodology employed throughout the paper is generally sound. That said, I do have a final set of comments to soften some of the causal claims made in the parole data study.

Statistical Analysis, Uncertainties, & Conclusions

The analyses and statistical reporting are generally properly conducted (effect size reporting, confidence intervals, etc.), except for one remaining comment with respect to failed adjustments for multiple comparisons in new S4 (which are easy to fix).

Clarity and Context

I appreciate that the authors followed my recommendations to make the paper converge faster to the key prediction and main finding and avoid the previous more piecemeal development of ideas. The fronted is much more focused now.

Remaining Comments

- Remove Causal Claims & Avoid Overgeneralizations: Make sure to avoid any overgeneralizations and causal claims in some of the correlational data. The claim that audiovisual glitches cause significant real-world inequality particularly in the parole study is not supported by actual causal evidence. While the parole study is definitely compelling as a general context, the observational nature of the archival data prevents any actual causal inference. For instance, glitch occurrence may still correlate with various institutional factors (e.g., time of day, courtroom tech quality, etc.) that might also correlate with parole outcomes. You might say that video glitches are associated with an amplification of inequality but you should temper any causal claims, especially in this study.

- Multiple Comparisons & Error Adjustment in S4: The authors test ten different glitch subtypes but I fail to see an error adjustment. Given the number of possible comparisons, you either need to fully transport all contrasts or indicate that you've used planned contrasts when collapsing and contrasting the ten conditions relative to control. Otherwise, the results suffer from an inflated alpha error. Instead of reporting all multiple comparisons (e.g., Bonferroni, Tukey etc.), it would be fine to report the main contrast but then report in the supplementary material which contrasts (after controlling for multiple comparisons) did vs. did not reach statistical significance.

Referee #4

(Remarks to the Author)

I thought this draft was substantially improved – substantially better. The motivation was beefed up a bunch (including the title change), and I liked the focus on uncanniness as an explanator, and ruling out other drivers alongside information and being prepared. I may have not fully remembered the details of the last version – in which case I am sorry – but the parole and some of the RCTs seem new (or more impactfully explained). Either way, I see this as a big improvement, alongside the better motivation and succinct policy section at the end. I happy to sign-off on this, so I'm a fully 180 conversion from reject to accept (so that was a very effective revision).

One piece of advice (request) from me would be to reduce the number of statistics in the text. I know Nature has a house style, but not sure how flexible that is, particularly for a paper with a lot of numbers in the text. It made it slow and hard to read, and I think could impact citations and reads. The best solution would be to have Nature relax a bit on this so not every number needs every statistic next to it every time. If that's not possible I might actually suggest citing less figures – so being more imprecise – to make the text easier to read.

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

Revision Notes for Ms. No. 2024-08-16049

Note: Our submission includes two manuscript versions, one with changes from the last round tracked, and one without track changes.

We would like to thank the editor and the review team for the opportunity to revise our manuscript, “Minor Audiovisual Glitches Undermine Virtual Communication” (Ms. No. 2024-08-16049), now titled “Video Call Glitches Trigger Uncanniness and Harm Consequential Life Outcomes.” We are grateful for the constructive comments and insightful suggestions made by the reviewers, and we appreciate the time and effort the editor invested in identifying the major areas for improvement: (1) the practical utility of our findings and (2) the technical robustness of our empirical approach. We have worked to address these concerns comprehensively, as detailed in our editor's letter and individual reviewer responses, resulting in what we believe is a significantly strengthened manuscript.

Underscoring Practical Significance

Our findings demonstrate practical significance through three converging lines of evidence: (1) meaningful impact on consequential decisions in the real world, (2) conservative individual-level measures that underestimate real-world comparative effects, and (3) cumulative disadvantages that amplify existing inequalities. We discuss each below:

First, **when audiovisual glitches can alter life-changing decisions, even modest effect sizes carry practical significance.** This is most evident in our parole hearings data. For parole hearings where an incarcerated individual glitched, the percent granted parole was 12% lower: 48.1% of hearings with glitches resulted in parole, compared to 60.1% without glitches. The affected individuals would undoubtedly view this systematic bias as profoundly consequential. Importantly, virtual parole hearings are still prevalent today—numerous states have permanently integrated them due to reduced costs, enhanced security, and improved public access¹⁻³—suggesting that glitches continue to pose a threat to the fairness of digital justice. Further, parole hearings represent just one domain within a broader landscape of high-stakes digital communication. Additional legal proceedings are migrating online⁴⁻⁶, including those dealing with traffic violations (which disproportionately burden low-income populations), custody determinations, and workers' compensation cases. Healthcare appointments and employment interviews have similarly shifted to virtual formats, with telehealth now offered in 79% of hospitals⁷, and video-based first-round interviews now comprising 80% of *all* initial job screenings⁸. Across these critical domains—legal, medical, and professional—we argue that the people affected would view any systematic disadvantage imposed by technical failures as fundamentally unjust. We have expanded the manuscript's discussion of high-stakes videoconferencing applications throughout (e.g., pages 9, 11, 21) to underscore this practical relevance and we summarize this discussion using a new chart (Extended Data Figure 1).

Second, the effect sizes we report often represent lower bounds on real-world impact, as **we used conservative, isolated measures (i.e., evaluating one person only), whereas real interpersonal evaluations tend to occur in comparative contexts.** For example, people routinely evaluate multiple candidates simultaneously—whether interviewing job applicants,

selecting healthcare providers, or choosing service professionals—which can amplify the relative influence of otherwise small differences between candidates. For example, if two similarly qualified candidates are assessed and one glitches, the resulting decrease in interpersonal evaluation can cost the affected individual a job opportunity.

To demonstrate this principle empirically, we include a new Study 3 where, rather than indicating trust in a health professional's recommendation (i.e., an evaluation in isolation, as in our prior telehealth study, now Study 2), we present participants with the choice of working with this health professional in person versus talking to another professional. When the health professional glitched, the percent of people opting to meet with them in person decreased by 16 percentage points (people choose to work with professional 61% of the time in the glitch condition, compared to 77% in the no-glitch control). This suggests that, even though Study 2 participants rated trust in the medical professional as relatively high in both conditions ($M_{\text{no-glitch}} = 5.05$, $M_{\text{glitch}} = 4.76$; measured in isolation on a 1-7 scale), this difference becomes much more consequential when put in a comparative context. In a new study, we find that simply having these minor glitches decreases a health professional's relative chance of obtaining a new client by 16 percentage points.

Finally, worth noting, we investigate the effects of glitches for one video call, but people participate in numerous video calls annually across work, healthcare, legal, and other domains. Thus, **seemingly discrete disadvantages can accumulate into substantial lifetime impacts**. This is particularly concerning when considered alongside documented disparities in internet access based on socioeconomic status, geography, housing status, age, and race⁹⁻¹³. Consequently, certain populations disproportionately experience glitches across multiple life domains, and these accumulated effects meaningfully diminish overall quality of life and risk amplifying existing social inequalities. In the revised manuscript, we added a concise summary of these citations (page 21).

Given our evidence that people underestimate the negative effects of glitches, we believe that documenting the extent to which glitches undermine high-stakes communication outcomes is worthwhile. Our research calls for policies that reduce the prevalence of glitches. For example, policymakers could bridge the digital divide by granting poorer communities high-speed internet access or offering third spaces for people without reliable internet access to use for consequential meetings. Further, when entities control the internet infrastructure that all parties must use—such as incarcerated individuals relying on the court internet for parole hearings—they must be held responsible for offering adequate internet. In 2016, the UN recognized the internet's fundamental role in enabling people to exercise their rights^{14,15}. However, for this recognition to have meaning, internet access must be reliable—poor connectivity arguably denies people the very rights the UN sought to protect. We now expound on this at the end of the paper (page 22).

Expanding Robustness of Results

Our revision now documents the robustness of our results to other types of glitches and provides additional evidence for our proposed process.

In the new Study 4, a comprehensive multi-wave experiment (N = 4945) demonstrates that several audiovisual glitch subtypes undermine interpersonal judgments. Further, we show that uncanniness captures predictive patterns in judgements that disruptiveness and magnitude measures miss. Specifically, after developing stimuli for the ten different glitch subtypes derived in Study 1 and validating their realism (Wave 1), we observe the uncanniness of each subtype (Wave 2) and measure the extent to which it predicts communicator evaluation (Wave 3). We conducted this study with three distinct samples following R1's helpful recommendation. This set-up allows us to avoid demand characteristics and prevent shared method variance (where both predictor and outcome measures are collected with the same measurement tools and the same participants, potentially inflating correlations).

Below, we detail how this multi-wave design addresses the core theoretical and methodological questions raised in the review:

- First, we addressed reviewer concerns about the realism of our stimuli: In Wave 1, we develop stimuli for ten different glitch subtypes (reflecting the typology derived in Study 1) and validate the realism and believability of these stimuli.
- In subsequent waves, we expand our inquiry beyond fleeting freezes and demonstrate the robustness of our results to other types of glitches: Eight out of the ten glitch subtypes significantly undermined interpersonal judgments (versus a no glitch control).
- We also address reviewers' calls for more robust mechanism evidence by leveraging variance in uncanniness by glitch subtype: We find that glitches vary in the extent to which they evoke uncanniness, and these feelings of uncanniness (in Wave 2) predict downstream interpersonal judgments (in Wave 3).

We also respond to the review team's questions about frustration. We agree with the team's intuitions that there is some form of negative affect (whether frustration or some other emotion) occurring due to glitches, which carries over and negatively impacts communication outcomes. Indeed, we are all too familiar with the frustration that arises when telehealth calls or online office hours go awry! After considering the review team's comments about frustration, we now treat it as part of our uncanniness causal chain, rather than a competing prediction.

At the end of the introduction, we re-orient alternative explanations under the overarching question: **why are glitches so frustrating?** We propose uncanniness as a key factor underscoring this reaction, but other plausible reasons exist. One is mere disruption—"this isn't going smoothly;" another is attribution—"you didn't prepare for our meeting;" a final one is information loss—"I can't understand what you are saying."

While these likely have some impact (which we consider and test throughout), we present a more ingrained and specific process of uncanniness that explains the effect of glitches more comprehensively. Combining the findings we presented in the last round with a new multi-wave study, we find that evidence that disruption and information loss can damage communication outcomes to some extent, but that uncanniness functions above and beyond these other processes and can more precisely explain when and why glitches are so costly. We expound on these

findings in a new section called Alternative Accounts (page 19). Here is a summary of relevant points:

Mere Disruption: Even when controlling for both objective and subjective disruptiveness, uncanniness still predicts negative hiring outcomes. For example, sustained freezes last twice as long as fleeting freezes and are perceived as more disruptive; yet, they're actually less uncanny and less damaging to hiring interest. Similarly, fleeting video loss is highly disruptive but low in uncanniness—and produces only small decreases in hiring interest, tracking with uncanniness rather than disruption level. These results demonstrate that uncanniness and disruptiveness are distinct constructs, and that uncanniness predicts evaluation patterns in a way that disruptiveness alone does not.

Information Loss. The core relationship between glitches and hiring holds even after controlling for information loss, bolstering our existing findings that glitches have negative effects even when they are designed to preserve information transfer. However, of note, audio-loss glitches had an outsized negative effect on hiring (vs. what their uncanniness scores alone would predict), suggesting information loss might be a parallel and additive process. In other words, audio-loss glitches undermine outcomes through *both* (1) uncanniness and (2) comprehension deficits.

Attribution. We chose not to consider attribution in this study (i.e., do you blame the communicator for the glitches) because we did not have a clear prediction for why glitch subtypes would affect attributions (other than due to magnitude). However, in other studies, we test whether participants attribute the glitch to the communicator or the third-party survey platform. The results are not affected by this attribution (see Supplement H).

We are grateful for the review team's suggestion to reconsider the role of frustration. We believe this reconceptualization (that uncanniness can explain *why* glitches are frustrating) better captures the psychological process driving the negative impact of glitches. An ancillary benefit of this reconceptualization is that our introduction and studies are tighter, as we now present a consolidated discussion of alternatives at the end of the paper.

Summary

In summary, this revision includes two all-new studies, one to address the concern of practical relevance (study 3), and one to address the concern of broader empirical robustness (study 4). We also expanded our discussion of practical relevance to high-stakes domains and underserved communities and reconceptualized frustration as part of our uncanniness causal chain rather than a competing explanation.

In doing so, we now include three new figures visualizing new data, two in the main manuscript and one in extended data. We also include one new infographic illustrating the prevalence of the video calling contexts, and finally, one table inventorying all glitch stimuli and placements (though we note that this last table is meant only for internal review purposes, as the human actors have not offered their permission to publish their likenesses). Lastly, we re-ordered the empirical package to improve flow and clarity:

	Original Manuscript	Revised Manuscript
1	Archival “get-to-know-you” dataset	<i>No change</i>
2	Hiring study [<i>moved to Supplement</i>]	Live Telehealth Webinar [<i>re-positioned</i>]
3	Live Telehealth Webinar	New: Telehealth Webinar (with mediation by uncanniness and choice outcome)
		Financial Advisor (with mediation) [<i>re-positioned</i>]
4	Financial Advisor (with mediation)	New: Multi-Wave Multi-Glitch Study
5	Face-to-Face versus Shared-Screen Study	<i>No change</i>
6	Archival Parole dataset	<i>No change</i>

Thank you again. Below, we respond to each reviewer. For ease of viewing, we provide the original reviewer comment in black, and our response in blue.

Reviewer 1:

Note: Our submission includes two manuscript versions, one with changes from the last round tracked, and one without track changes.

Thank you very much for your thoughtful and constructive review of our manuscript. We are glad the work's novelty and importance came through, and your feedback was invaluable in strengthening our paper. We have carefully addressed each of your comments in the revised manuscript, with details provided below:

1. “While the effect documented in the paper has clear relevance for individuals who carry out important meetings remotely (e.g., job interviews), the paper makes a societal claim of substantive importance: the effect of glitches risks entrenching inequalities in society. I am positively predisposed towards this claim, and this is a comment aimed at sharpening this aspect of the paper. A devil’s advocate could argue that by and large people across the SES ladder rely on the same data networks nowadays. Is it really true that glitches are more frequently experienced by low SES citizens? For example, to the extent that data connectivity is better in urban centers and that the poor tend to concentrate in urban areas (vs. less densely populated areas surrounding cities), one might even make the opposite prediction. I am not writing this because I believe your contention to be false but only to illustrate the need for validating a key assumption you rely on to claim that glitches exacerbate inequality.”

Thank you for encouraging us to provide stronger support for our claim that glitches entrench societal inequalities. In our revision, we now support our assertion that glitches are more frequently experienced by low-SES citizens by citing several papers that empirically document inequality in internet access among lower-income neighborhoods^{10,12,13}. In the paper, we now also point to other underserved communities that disproportionately suffer from lower-quality internet, such as rural areas and senior citizens^{9,11}. These additional citations support our assertion that glitches—which are amplified in the context of poor internet—exacerbate inequalities, emphasizing the societal relevance of our work.

2. “I think that the paper is currently a bit hand-wavy on the psychological mechanism. You rely on the construct of “uncanniness” but I struggle to fully understand the nature and outlines of this process. You define it as the reaction accompanying a violation of expectations. Could one argue that this construct is too close to the independent variable? (More akin to a manipulation check than a psychological process?) I think the idea makes intuitive sense but perhaps the paper could conceptualize this construct with greater precision. I’ll note that despite the large amount of data, the paper contains relatively little process evidence by moderation (in this sense I found Study 5 very important).”

Thank you for encouraging us to add more clarity regarding the psychological mechanism. We now conceptualize our process with more precision in the upfront, theoretical section of the paper. We now borrow extensively from prior work in technology’s ability to convey high levels of social presence¹⁶⁻¹⁸, as well as theories surrounding uncanniness as a so-called “break” in presence¹⁹⁻²¹.

Specifically, we now begin by arguing that video calling offers a high level of social presence relative to other communication technologies, based on its audiovisual richness²² realism and synchrony²³. We borrow from the conceptualization of high social presence as an “illusion of non-mediation,” where the technological medium becomes invisible and users begin to behave as if it were real life²⁴.

We next theorize that glitches break this illusion of face-to-face contact afforded by video calling, as they do not occur in real life, and this break triggers uncanniness. This prediction builds on prior work looking at artificial agents like robots or videogame animations, in which scholars proposed that uncanniness occurs after a “break in presence,” wherein observers must abruptly recategorize an agent from “human” to “virtual.”

Thank you for encouraging us to tighten our theoretical support. We believe that this updated theory highlights our conceptual contribution by clearly bridging previously siloed work in the social presence of technology and the uncanny breaks in presence triggered by artificial agents.

Further, with this updated conceptualization, **we are better able to disentangle uncanniness from the independent variable (i.e., glitchiness) both conceptually and empirically:**

- We argue that glitches (the independent variable) are events that occur during video calls, whereas uncanniness (the mechanism) is a psychological reaction to that event or events. We demonstrate this empirically through traditional tests of discriminant validity: In all studies that measure both the manipulation check and the mechanism, (a) the 95% confidence intervals surrounding their correlation do not include 1 (considered evidence of discriminant validity)²⁵, and (b) exploratory factor analyses yield two separate factors, one for each multi-item construct measure (See Supplement E for results).
- Supporting this, we find that not all glitches cause uncanniness. We show this in our previous submission by manipulating whether or not a break in presence occurred. We investigated glitches in contexts with lower (i.e., shared screen presentation) versus higher (i.e., face-to-face video call) social presence (Study 5).
- Further, the extent of glitchiness (as measured by our manipulation check) is not a perfect predictor of uncanniness. We show this in a new study by introducing multiple different glitch types, which allows for variance in both how “glitchy” and how uncanny different glitches are. We find that some glitch subtypes are perceived as relatively severe, but not particularly uncanny. To illustrate this more concretely, we present one case study, comparing fleeting freezes (<2 seconds) and extended freezes (4 seconds). Despite extended freezes seeming significantly more disruptive (vs. fleeting freezes), they were significantly less uncanny and had a smaller negative effect on interpersonal judgment.

This is supported by prior work finding that factors beyond merely size of distortion play a role in determining what feels natural versus unnatural. One example that was illuminating for us is Schein and Gray (2015)²⁶, which finds that an eye-less person is uncanny but a nose-less person is not (vs. a non-distorted image), even though these distortions are similar in

magnitude. As understanding the aesthetic drivers of uncanniness is highly complex—indeed, Schein and Gray (2015)²⁶ was dedicated to identifying *one* specific (and static!) factor driving uncanniness—we do not attempt to understand all the multi-modal and dynamic features of different glitches that elicit uncanniness. Rather, we use this variation to demonstrate uncanniness’s empirical value for understanding why glitches undermine interpersonal judgment (as you suggested below).

3. “Related to the previous, and as you also note in the paper, it is relatively easy to argue for alternative psychological processes. For example, lower rates of parole in virtual court appearances with glitches may be the result of the judge’s frustration. Glitches in video calls are frustrating indeed.”

Thank you for raising the question of frustration. We agree with you that there is some form of negative affect (whether frustration or some other emotion) occurring due to glitches, which carries over and negatively impacts communication outcomes. Indeed, we are all too familiar with the frustration that arises when telehealth calls or online office hours go awry! After reconsidering the role of frustration, we now treat it as part of our uncanniness causal chain, rather than a competing prediction.

Specifically, we now reorient alternative explanations under the overarching question: **why are glitches so frustrating?** We propose uncanniness as a key factor underscoring this reaction, but other plausible reasons exist. One is mere disruption—“this isn’t going smoothly;” another is attribution—“you didn’t prepare for our meeting;” a final one is information loss—“I can’t understand what you are saying.”

In our multi-wave study (inspired by your next point), we find evidence that mere disruption and information loss play a partial role in the negative impact of glitches. Importantly, though, uncanniness (1) affects outcomes even when accounting for these other processes and (2) explains the results in cases where these alternatives do not hold. Here are a few findings illustrating these points (which can be found in a new section called Alternative Accounts, page 19):

Mere Disruption: Even when controlling for both objective and subjective disruptiveness, uncanniness still predicts negative hiring outcomes. One example mentioned above: sustained freezes last twice as long as fleeting freezes and are perceived as more disruptive; yet, they’re actually less uncanny and less damaging to hiring interest. Similarly, fleeting video loss is highly disruptive but low in uncanniness—and produces only small decreases in hiring interest, tracking with uncanniness rather than disruption level. In short, uncanniness predicts evaluation patterns in a way that mere disruptiveness does not.

Information Loss. The core relationship between glitches and hiring holds even after controlling for information loss, bolstering our existing findings that glitches have negative effects even when they are designed to preserve information transfer. However, of note, audio-loss glitches had an outsized negative effect on hiring (vs. what their uncanniness scores alone would predict), suggesting information loss might be a parallel and additive process. In other words, audio-loss glitches undermine outcomes through *both* (1) uncanniness and (2) comprehension deficits.

We have included a new Extended Data Table 3 that helps to illustrate our points: Setting aside the two audio-loss subtypes, the hiring variable shows a similar gradient pattern (increasing color saturation) to uncanniness, relative to comprehension and disruption (which show more diffuse color saturation).

Glitch Type	Uncanniness <i>Red = more uncanny vs. control</i>	Comprehension <i>Red = lower comprehension vs. control</i>	Disruption <i>Red = more disruptive vs. control</i>	Hiring Interest <i>Red = larger drop vs. control</i>
Sustained Freeze	0.52	−0.03	3.27	−0.21
Distorted Video	0.59	−0.03	3.03	−0.16
Fleeting Video Loss	0.68	−0.02	3.33	−0.28
Lag	0.74	0.00	2.02	−0.12
Static	0.89	−0.05	2.58	−0.53
Fleeting Audio Loss	1.11	−0.18	3.99	−1.27
Sustained Audio Loss	1.20	−0.23	4.22	−1.55
Echo	1.38	−0.05	2.90	−0.52
Fleeting Freeze	1.41	0.02	3.00	−0.66
Sustained Video Loss	1.78	−0.05	4.04	−0.49

4. “You treat glitches as “one thing” but there are many types of glitches, and some may be more strongly associated to feelings of uncanniness and to negative interpersonal outcomes. For example, I find lack of synchrony between sound and image very uncanny. Even if other types of glitches may be objectively involving a much larger information loss, if your hypothesis is correct they may not produce effects that are as negative. Going back to the point above about moderation, one option could be to build a typology of glitches, pretest them for uncanniness, and then conduct a process test via moderation. I like the classification provided for Study 1 (p 26) and perhaps more could be done in this direction. I am mentioning this as an opportunity to deepen your insights and implications. I do not think that the paper needs necessarily more studies given the amount of data that you already gathered.”

We appreciate your encouragement to explore the broader glitch landscape beyond fleeting freezes (which we originally focused on to minimize confounds with information loss). We now include a comprehensive, multi-wave study (Study 4) that directly implements your suggested design.

After developing stimuli for the ten different glitch subtypes derived in Study 1 and validating their realism (Wave 1), we observe the uncanniness of each subtype (Wave 2) and measure the extent to which it predicts communicator evaluation (Wave 3). This new multi-wave experiment (N = 4945) demonstrates that several audiovisual glitch subtypes undermine interpersonal judgments, and that uncanniness captured predictive patterns in judgements that disruptiveness or magnitude measures missed.

Thank you so much for suggesting this enormously helpful design-- as you can see, we have already mentioned it multiple times in this letter! We hope you agree that this new study increases the ecological validity and robustness of our work while also underscoring the predictive value of uncanniness.

5. “A smaller comment: could the effect reverse as a function of the topic of conversation? It is a comedy cliché for someone on a call to pretend that the connection is poor when facing scrutiny or a difficult conversation. Could it be that glitches that hurt the ability to communicate may sometimes give “cover” to people who find themselves in a difficult spot?”

Others share your logic: The tool we used to add glitches into a live Zoom webinar in Study 2 (formerly Study 3) was originally designed to help people exit calls! After considering your comment, competing predictions emerge: On one hand, glitches could be helpful in difficult conversations. But on the other hand, negative conversation topics could exacerbate glitches’ negative effects due to “stacking” of discomfort. While we did not manipulate valence in our experiments in this project, we think it is an excellent potential future avenue for research.

Thank you again for your constructive comments and suggestions. Your feedback was extremely helpful in strengthening our manuscript, and we truly appreciate the time and thought you put into your review.

Reviewer 2:

Note: Our submission includes two manuscript versions, one with changes from the last round tracked, and one without track changes.

We very much appreciate your comments and feedback. We are pleased to hear your positive appraisal of our robust corpus of studies and real-world application. Below, we quote your concerns and provide our responses:

1. “This is a robust corpus of studies, they have good power, applied across a variety of contexts that show a large spectrum of real world applications. My overarching concern stems from how practically significant these effects are. Let's say the mean difference in outcome with glitches is 4.25 vs 4.5 and the effect size is significant. But is this a practically meaningful difference. In other words, is 4.25 actually low enough to be concerned about glitches. Yes glitches, do affect the intended outcome somewhat adversely, but it appears the intended outcome is still being achieved.”

We appreciate this comment about practical significance, which is a crucial consideration beyond statistical significance alone.

We want to clarify that our Cohen's D s (which can be viewed in Table 1)—considered “medium” effect sizes according to common standards²⁷—are conservative estimates of our effect. In our studies, **we used isolated measures (i.e., evaluating one person only), whereas real interpersonal evaluations tend to occur in comparative contexts.** For example, people routinely evaluate multiple candidates simultaneously—whether interviewing job applicants, selecting healthcare providers, or choosing service professionals—which can amplify the relative influence of otherwise small differences between candidates. For example, if two similarly qualified candidates are assessed and one experiences technical difficulties, the resulting decrease in interpersonal evaluation can cost the affected individual a job opportunity.

We conducted a new study to demonstrate this. In the new Study 3, rather than indicating trust in a health professional's recommendation (i.e., an evaluation in isolation, as in our prior telehealth study, now Study 2), we present participants with the choice of working with this health professional in person versus talking to another professional. When the health professional glitched, the percent of people opting to meet with them in person decreased by 16 percentage points (people choose to work with professional 61% of the time in the glitch condition, compared to 77% in the no-glitch control). This suggests that, even though Study 2 found that participants in both conditions rated trust in the medical professional (measured in isolation, on a 1-7 scale) as relatively high ($M_{\text{no-glitch}} = 5.05$, $M_{\text{glitch}} = 4.76$), this difference becomes much more consequential when put in a comparative context. In a new study, we find that simply having these minor glitches decreases a health professional's relative chance of obtaining a new client by 16 percentage points.

Beyond effect size magnitude, we believe two additional factors demonstrate the real-world importance of our findings: (1) the consequential outcomes we examine and (2) the potential for cumulative harm.

Consequential Outcomes: When audiovisual glitches can alter life-changing decisions, even modest effect sizes carry practical significance. This is most evident in our parole hearings data. For parole hearings where an incarcerated individual glitched, the percent granted parole was 12 percentage points lower: 48.1% of hearings with glitches resulted in parole, compared to 60.1% without glitches. Further, parole hearings represent just one domain within a broader landscape of high-stakes digital communication: Healthcare appointments and employment interviews have similarly shifted to virtual formats, with telehealth now offered in 79% of hospitals⁷, and video-based first-round interviews now comprising 80% of *all* initial job screenings⁸. Across these critical domains—legal, medical, and professional—we argue that the people affected would view any systematic disadvantage imposed by technical failures as profoundly consequential and unjust.

Cumulative Impact: Videoconferencing has permeated nearly all forms of interaction, spanning professional meetings, healthcare appointments, job interviews, and legal consultations and proceedings. Thus, seemingly discrete disadvantages can compound over time. This is particularly concerning when considered alongside documented disparities in internet access based on socioeconomic status, geography, housing status, age, and race⁹⁻¹³. Consequently, certain populations disproportionately experience glitches across multiple life domains, and these accumulated effects meaningfully diminish overall quality of life and risk amplifying existing social inequalities. For example, the same individual may participate in dozens of virtual interactions annually across work, healthcare, and legal domains, with even minor glitch-related disadvantages in each setting compounding into substantial cumulative effects compared to an individual who rarely experiences glitches. We have expanded this discussion in the manuscript (page 21) to provide additional support for these claims.

2. “The study maybe statistically robust and exhaustive, but it fails to drive the importance of the questions and the context. And glitches ultimately will reduce as technologies and access both improve with reduced costs.”

This comment addresses a fundamental assumption underlying our research: If videoconferencing glitches are a temporary problem, our results would provide limited value. However, we believe the impact of glitches will persist as a significant phenomenon for several reasons:

First, new technologies introduce new failure sources even as they solve old ones. Consider the evolution of communication technology: while phone calls may experience fewer technical failures today, videoconferencing brought about new glitches. This pattern continues with next-generation technologies: Zoom's recent introduction of 3D group meetings²⁸ demands significantly more bandwidth and processing power than standard 2D meetings, again increasing vulnerabilities to glitches. This represents a fundamental characteristic of technological evolution: as infrastructure capacity expands, we simultaneously develop more demanding applications that strain these very improvements. Glitches do not disappear over time—they migrate, following technology as it evolves into new formats.

Moreover, as communication technologies increasingly simulate real-life interactions, our theoretical framework predicts that glitches may become even *more* uncanny. That is, as technology's capacity for creating social presence increases, so too does the risk for uncanniness. Consequently, advances in communication technology may *intensify* rather than reduce the impact of technical glitches, suggesting that our research may gain even greater relevance as communication technologies evolve. Other scholars have proposed complementary perspectives: as time goes on, people become increasingly discerning with technology, which causes uncanniness to persist²⁹.

Lastly, we find it worth noting, even if glitches do diminish in platforms like Zoom, this improvement will likely be uneven—technical problems are much more likely to disappear first for the “haves” than the “have-nots.” This disparity represents an important consideration for policymakers addressing digital equity given the impact glitches can have in consequential domains.

In the revised manuscript, we have added content (page 24) to better support our claims and make the points you raise more salient to the reader. Thank you very much.

Reviewer 3:

Note: Our submission includes two manuscript versions, one with changes from the last round tracked, and one without track changes.

We sincerely appreciate your positive evaluation and the care you took in suggesting ways to improve both our methodological and conceptual rigor. Below, we address the areas for improvement you identified to further strengthen our claims and contribution.

1. “More Precise Theory Foundation: The empirical package is solid and the finding is interesting and important. What the authors could carve out more assertively is the precise causal chain and theoretical foundation to link (1) the technical properties of glitches to (2) a distinct psychological process triggered by these disruptions (why uncanniness and not other mechanisms?), and (3) the specific downstream consequence on interpersonal judgments and decisions. It would be helpful to formulate somewhat sharper conceptual boundaries to differentiate why glitches specifically trigger uncanniness rather than more immediate responses like frustration or competence attributions. I also felt that the authors could situate the theory-foundation more in classic theories of social presence or communication synchrony. This would help explain more precisely the psychological pathway from technical disruption to social judgment and would in turn also help to carve out the specific contribution in light of classic theories of communication synchrony or social presence.”

Thank you for encouraging us to add clarity and precision to our theoretical foundation.

We have carefully considered each of the three causal-chain elements you highlight, as well as why uncanniness is more central to the predictions than other mechanisms like frustration. In the order you highlighted them...

(1) Technical properties of glitches:

We agree that it would be interesting to link specific technical properties of glitches to uncanniness and interpersonal judgment. However, we feel this area is outside the scope of the present work. In fact, although research on non-human agents finds variance in the extent to which uncanniness is produced by various design elements, it does not yet offer a parsimonious framework for understanding this variance, likely due to the complexity of dynamic multi-sensory processing. In fact, some papers are dedicated to uncovering the role of a *single* aesthetic feature in uncanniness perception in static images (e.g., the eyes²⁶). However, we think this is a fascinating avenue for future research and we feature it on page 23.

With that said, we did expand our inquiry to 10 different glitch subtypes (vs. the single type, fleeting freeze, in the prior version of the manuscript) to demonstrate the robustness of our effect across glitches with different technical properties. Further, we measured the relative uncanniness of each glitch subtype and examined the relationship between uncanniness and interpersonal judgments. This leads us to the subsequent element in the causal chain—uncanniness—which we cover next.

(2) Distinct psychological process of uncanniness and not others: We will answer this in two parts.

First, to carve out the precise causal chain from the presence of glitches to uncanniness, we followed your helpful recommendation to draw upon social presence and media synchrony theories. Specifically, we now begin by arguing that video calling offers a high level of social presence relative to other communication technologies, based on its audiovisual richness²² realism and synchrony^{23,24}. We borrow from the conceptualization of high social presence as an “illusion of non-mediation,” where the technological medium becomes invisible and users begin to behave as if it were real life²⁴.

We next theorize that glitches break this illusion of face-to-face contact afforded by video calling, as they do not occur in real life, and this break triggers uncanniness. This prediction builds on prior work looking at artificial agents like robots or videogame animations, in which scholars proposed that uncanniness occurs after a “break in presence,” wherein observers must abruptly recategorize an agent from “human” to “virtual.”

Thank you for your suggestion to incorporate social presence and media synchrony theories into our conceptual development. As you can see, these served as pillars in our reconceptualization.

Second, after reflecting on the review team’s comments about frustration, we now treat it as part of our uncanniness causal chain, rather than a competing prediction. Specifically, we now reorient alternative explanations under the overarching question: **why are glitches so frustrating?** We propose uncanniness as a key factor underscoring this reaction, but other plausible reasons exist. One is mere disruption—“this isn’t going smoothly;” another is attribution—“you didn’t prepare for our meeting;” a final one is information loss—“I can’t understand what you are saying.

We add new support for our proposed process in our multi-wave study: although we find evidence that disruptiveness and information loss play a partial role in the negative impact of glitches, uncanniness (1) affects outcomes even when accounting for these other processes and (2) explains the results in cases where these alternatives do not hold. Here are a few findings illustrating these points (which are expounded upon in greater detail in a section called Alternative Accounts, page 19):

Mere Disruption: Even when controlling for both objective and subjective disruptiveness, uncanniness still predicts negative hiring outcomes. One example mentioned above: sustained freezes last twice as long as fleeting freezes and are perceived as more disruptive; yet, they’re actually less uncanny and less damaging to hiring interest. Similarly, fleeting video loss is highly disruptive but low in uncanniness—and produces only small decreases in hiring interest, tracking with uncanniness rather than disruption level. In short, uncanniness predicts evaluation patterns in a way that mere disruptiveness does not.

Information Loss. The core relationship between glitches and hiring holds even after controlling for information loss, bolstering our existing findings that glitches have negative effects even when they are designed to preserve information transfer. However, of note, audio-loss glitches

had an outsized negative effect on hiring (vs. what their uncanniness scores alone would predict), suggesting information loss might be a parallel and additive process. In other words, audio-loss glitches undermine outcomes through *both* (1) uncanniness and (2) comprehension deficits.

Attribution. We chose not to consider attribution in this study (i.e., do you blame the communicator for the glitches) because we did not have a clear prediction for why glitch subtypes would affect attributions (other than due to magnitude). However, in other studies, we test whether participants attribute the glitch to the communicator or the third-party survey platform. The results are not affected by this attribution (see Supplement H).

We have included a new Extended Data Table 3 that helps to illustrate our points: Setting aside the two audio-loss subtypes, the hiring variable shows a similar gradient pattern (increasing color saturation) to uncanniness, relative to comprehension and disruption (which show more diffuse color saturation).


Glitch Type	Uncanniness <i>Red = more uncanny vs. control</i>	Comprehension <i>Red = lower comprehension vs. control</i>	Disruption <i>Red = more disruptive vs. control</i>	Hiring Interest <i>Red = larger drop vs. control</i>
Sustained Freeze	0.52	−0.03	3.27	−0.21
Distorted Video	0.59	−0.03	3.03	−0.16
Fleeting Video Loss	0.68	−0.02	3.33	−0.28
Lag	0.74	0.00	2.02	−0.12
Static	0.89	−0.05	2.58	−0.53
Fleeting Audio Loss	1.11	−0.18	3.99	−1.27
Sustained Audio Loss	1.20	−0.23	4.22	−1.55
Echo	1.38	−0.05	2.90	−0.52
Fleeting Freeze	1.41	0.02	3.00	−0.66
Sustained Video Loss	1.78	−0.05	4.04	−0.49

Interestingly, you had characterized frustration as an “immediate reaction.” In contrast, you can see that we now position frustration as an *outcome* that emerges from uncanniness and these aforementioned processes, with our dependent variables capturing meaningful consequences of negative affect. Our hope is that we now add nuance and richness to this understanding by going upstream to examine what causes this negative affect in the first place—specifically, how the uncanniness triggered by technical glitches contributes to these frustration responses.

4) Specific downstream consequence

As alluded in our response above, we conceptualize the outcome measures as meaningful consequences of negative affect—in other words, a negative halo. This dovetails with other findings documenting negative downstream consequences of encounters with the uncanny^{21,30}. To that end, we now present the “B link” of our theoretical model as a spillover effect, whereby the negative experience of uncanniness contaminates subsequent judgments and evaluations (page 4).

We include a diagram below that represents our conceptual model, including a red, bold line for our primary mechanism (uncanniness), solid lines for consistent but insufficient mechanism evidence (mere disruption), and dashed lines for inconsistent or weak mechanism evidence (information loss and competence attribution). We do not currently feature it in the manuscript, but we certainly can if the review team sees fit.



2. “Deeper Insight on Variation in Glitches & Additional Model Specifications: Looking at Table 2 raises the question whether the effect is truly uniform or if the authors could better exploit variation in glitches. Specifically, the effect sizes look much stronger for video compared to mere audio distortions, while more fleeting vs. more sustained disruptions tend to reduce effect sizes. While this variation could justify the focus on video-based interruptions in the experimental work, it also presents opportunities for a deeper empirical investigation. I think some additional analyses of Studies 1 and 6 would be particularly useful. In Study 1, the authors’ didn’t model interaction effects between different types of glitches (e.g., combined audio-visual glitches) or temporal sequences that are more vs. less disruptive. For example, I would intuitively expect that it makes a difference whether I experience many short disruptions or a few but long and sustained glitches. In Study 6, the authors could also use a LASSO with the glitch dictionary features as predictors and test specifically which types or sequences of glitches evoke the most negative impact on parole outcomes. Such an analysis has the potential to reveal whether certain types or combinations of glitches are particularly damaging in such high-stakes decisions. To be clear: I consider these additional analyses not absolutely critical for the paper’s core contribution, but they would deepen our understanding of potentially critical “glitch patterns” of when and how technical disruptions are most problematic.”

Thank you for encouraging us to better tease apart the variation in glitches and their potentially heterogeneous effects—this suggestion was at the foundation of our revision effort. Your suggested analyses for Studies 1 and 6 were extremely clever. Unfortunately, when we explored these possibilities, we determined that our studies lack sufficient power for the sophisticated analyses you proposed. In Study 1, each glitch subtype had an average of 35 observations, making interaction analyses unreliable. Additionally, while we share your curiosity about the timing or sequencing of glitches, this dataset contains an enormous amount of noise in people’s glitch experiences: some experienced multiple glitch types simultaneously or sequentially; others experienced the same glitch repeatedly. Thus, the resource-intensive process of watching over

100 hours of video to identify specific timestamps seemed unlikely to yield clear insights. In Study 6, with 472 observations with over 20 glitch dictionary features, we lack adequate power for LASSO regression with binary outcomes (which we understand to require 10 observations per outcome level [parole = yes vs. no] per dictionary feature).

However, we hope that some of the findings from our new multi-wave study may provide some interesting insights regarding glitch heterogeneity. A few highlights:

- Different glitches produce different levels of uncanniness (consistent with prior work), and this uncanniness predicts the negative outcomes on evaluations
- The extent of uncanniness is not easily explained through magnitude or duration of glitches. For instance, even though extended freezes were twice as long as fleeting freezes and were perceived as more disruptive, they were significantly less uncanny and had a smaller negative effect on interpersonal judgment.
- The glitches that directly compromise information transfer (i.e., the audio-loss glitches) have an outsized effect on hiring, beyond what uncanniness would predict. This suggests information loss is a parallel process alongside uncanniness. In other words, audio loss glitches undermine outcomes through both uncanniness and comprehension deficits.
- Beyond these observations, it is not readily apparent what drives differences in uncanniness. For example, the three **most** uncanny glitches (echo, fleeting freeze, and sustained video loss) and the three **least** uncanny glitches (sustained freeze, distorted video, fleeting video loss) all involve both audio and visual distortions, longer and shorter durations, as well as freezes and video loss. As we mentioned above, the existing literature does not provide a clear roadmap for understanding what technical features drive uncanniness. For this reason, we were reluctant to speculate based on this dataset. We hope this is the first of many investigations.

“Robustness of Glitch Manipulation: All experimental studies induce frequent and short video distortions (freezing the video while showing the person). This raises the question whether your experimental studies are sensitive or relatively robust to alternative glitch interventions? For example, I kept wondering whether you would produce the same effect using (1) a blackscreen or (2) a pixelated image, instead of the video freeze. If these alternative disruptions produce weaker effect, this would even strengthen the authors’ predicted theory and at the same time offer practical guidance on how to deal with it (e.g., immediately block the screen to avoid showing the target person).”

This comment was extremely useful in helping us design and conduct the new Study 4 (discussed above), which looks at the robustness of the glitch effect across multiple subtypes varying in duration and modality. We hope that our discussion in the point above satisfies your curiosities about how these new data contribute to the overall empirical package and further illustrate the predictive power of uncanniness.

We also appreciate and share your desire for practical guidance. Of note, we expected that minimizing uncanniness would prove challenging, as prior work has alluded to an uncanny ‘wall’ (vs. valley) that is “an impossible traverse”²⁹. The prior version of our manuscript already included 3 intervention attempts with unimpressive results (Supplements I and J).

Thus, we explored a few potential solutions meant to prevent uncanniness from being triggered in the first place. For example, we piloted a few conditions where communicators switch from video to audio-only prior to any glitches. Unfortunately, these strategies seemed to introduce new issues (e.g., disappointment, confusion). While we regret not being able to offer a stronger recommendation, our shared-screen study suggests that switching away from a face-to-face simulation (e.g., by sharing one’s screen) may be a promising workaround.

Glitches are a complicated problem that may be difficult to address individually. Therefore, rather than offering individual-level solutions, we now conclude the paper with a call for more systemic actions on the part of policymakers and governing bodies to increase internet equity and access.

“Focus & Positioning: As highlighted earlier, the manuscript is generally well-written. But the opening section would benefit from a more streamlined and precise articulation of the paper's core contribution. The authors should more directly establish that seemingly trivial technical glitches can significantly impact social judgments and consequential decisions across various contexts (hiring, healthcare, parole decisions, etc.). Rather than gradually building to this key main finding, I felt that the introduction could more immediately highlight this paradox and offer a more focused opening. This would not only boost reader engagement but would also carve out the paper's theoretical and practical significance much faster and more direct. The somewhat gradual and piecemeal development of ideas almost dilutes the broad importance of the paper.”

This was an immensely valuable comment for us in the revision process. In the revised manuscript, we have streamlined the introduction to provide the focus of inquiry more quickly, and we more strongly feature the broader relevance of our effects for social judgements in consequential decisions. In the prior manuscript, we arrived at our core finding on line 36 of the manuscript. Now, we arrive at the core finding on line 17.

Minor

4.5 “Code documentation & Data: Please revise your markdown files to (1) use relative instead of absolute file paths (to make it easier for people to replicate your results) and (2) make sure to use groundhog or other version control packages to make sure your code runs with future R and package updates (write code with a focus on better reproducibility). Finally, please also upload the data and code for S6.”

We have updated our R files in line with your reproducibility and replicability requests. Note that, rather than using groundhog (which was not integrating well with R Markdown), we have opted to provide a complete list of the exact package specifications used and how to install them (based on recommendations and language from Data Colada³¹).

Regarding Study 6, we have now posted the code. We have not posted the data due to this language in our IRB documentation: “Given the sensitive nature of the research and the population, I do not intend to share hearing-level data (even in de-identified form) on online open-science repositories (despite this being the increasing norm in social sciences research) in an abundance of caution. Researchers may request access directly through their IRB or Open Record Request.”

As you see, we can provide data directly to individual researchers, and we would be more than happy to send you the materials through a password-protected link.

Thank you once again for your contribution to our manuscript through your insightful comments. We believe making the changes you suggested has resulted in a much-improved manuscript.

Reviewer 4:

Note: Our submission includes two manuscript versions, one with changes from the last round tracked, and one without track changes.

Thank you for your comments and suggestions. Your comments were very helpful in enabling us to clarify our contribution. Below we quote your concerns and then summarize the steps we took to address each one:

“1) The results are not surprising. Anyone that has used Zoom, Teams, Webex or any other video conferencing technology would know that lower quality video calls are frustrating, and some of that frustration can potentially be transferred to the other partner in the call. The RCTs here involved considerable “glitches” – four or five pauses of up to 2 seconds in a 75 second call, which is a lot. This would involve about 10% of the video call time frozen. As such the results seems somewhat unsurprising in that a repeatedly freezing call renders lower post-call feedback scores. Indeed, I might suspect that, for example, if somebody presented on a Zoom call with very bad lighting, a background with many distractions (like kids fighting), or fuzzy bad audio could have similar impacts.”

To your first argument—that the results are not surprising—we find that not everyone would agree. The paper presents evidence that people do not intuit glitches’ negative impacts (see Extended Data Table 1). In our qualitative interviews with sales professionals *and* in our surveys, we find that people overwhelmingly believe that glitches are trivial. We see this as a rational expectation, as glitches are frequent and often outside our control.

To your second argument—that “lower quality video calls are frustrating, and some of that frustration can potentially be transferred to the other partner in the call”—we agree that there is negative affect (perhaps frustration) that occurs due to glitches, which carries over and negatively impacts communication outcomes. Upon deeply considering this comment (and those of the review team more broadly), we have revised our theorizing and now treat frustration as part of our uncanniness causal chain, rather than a competing prediction.

Specifically, we now reorient alternative explanations under the overarching question: **why are glitches so frustrating?** We propose uncanniness as a key factor underscoring this reaction, but other plausible reasons exist. One is mere disruption—“this isn’t going smoothly;” another is attribution—“you didn’t prepare for our meeting;” a final one is information loss—“I can’t understand what you are saying.” While these likely have some impact (which we consider and test throughout), we present a more ingrained and specific process of uncanniness that explains the effect of glitches more comprehensively.

We add new support for this process in our multi-wave study: although we find evidence that disruptiveness and information loss play a partial role in the negative impact of glitches, uncanniness (1) affects outcomes even when accounting for these other processes and (2) explains the results in cases where these alternatives do not hold. Here are a few findings

illustrating these points (which are expounded upon in a section called Alternative Accounts, page 19):

Mere Disruption: Even when controlling for both objective and subjective disruptiveness, uncanniness still predicts negative hiring outcomes. A key example: Sustained freezes last twice as long as fleeting freezes and are perceived as more disruptive; yet, they're actually less uncanny and less damaging to hiring interest. Similarly, fleeting video loss is highly disruptive but low in uncanniness—and produces only small decreases in hiring interest, tracking with uncanniness rather than disruption level. These results demonstrate that uncanniness and disruptiveness are distinct constructs, and that uncanniness predicts evaluation patterns in a way that disruptiveness alone does not.

Information Loss. The core relationship between glitches and hiring holds even after controlling for information loss, bolstering our existing findings that glitches have negative effects even when they are designed to preserve information transfer. However, of note, audio-loss glitches had an outsized negative effect on hiring (vs. what their uncanniness scores alone would predict), suggesting information loss might be a parallel and additive process. In other words, audio-loss glitches undermine outcomes through *both* (1) uncanniness and (2) comprehension deficits.

We have included a new Extended Data Table 3 that helps to illustrate our points: Setting aside the two audio-loss subtypes, the hiring variable shows a similar gradient pattern (increasing color saturation) to uncanniness, relative to comprehension and disruption (which show more diffuse color saturation).

Glitch Type	Uncanniness <i>Red = more uncanny vs. control</i>	Comprehension <i>Red = lower comprehension vs. control</i>	Disruption <i>Red = more disruptive vs. control</i>	Hiring Interest <i>Red = larger drop vs. control</i>
Sustained Freeze	0.52	−0.03	3.27	−0.21
Distorted Video	0.59	−0.03	3.03	−0.16
Fleeting Video Loss	0.68	−0.02	3.33	−0.28
Lag	0.74	0.00	2.02	−0.12
Static	0.89	−0.05	2.58	−0.53
Fleeting Audio Loss	1.11	−0.18	3.99	−1.27
Sustained Audio Loss	1.20	−0.23	4.22	−1.55
Echo	1.38	−0.05	2.90	−0.52
Fleeting Freeze	1.41	0.02	3.00	−0.66
Sustained Video Loss	1.78	−0.05	4.04	−0.49

Unlike these alternative explanations that predict uniform negative effects of glitches across contexts, our uncanniness theory makes a novel prediction: glitches should be most damaging in high social presence contexts. Study 5 confirms this—while glitches feel similarly disruptive during shared-screen presentations versus face-to-face video calls, they're less (more) uncanny and less harmful to interpersonal judgments in the lower (higher) social presence context.

“2) “The policy implications are limited. Yes, it is true that individuals with lower incomes likely have worse internet connections on average (although this correlation may not be high, as lower income individuals often live in high-density urban areas which have good connections). As such they may suffer somewhat more here. But the scale of this seems limited. By 2024 most low income individuals are working and living fully in person, with little zoom activity. The people that tend to have remote or hybrid jobs are higher income professionals, who usually have reasonable internet connection (as this is required for their jobs). So, yes, in the paper they had a nice example from a higher stakes law example, but this was during the pandemic. These kinds of high stakes interactions – like law courts – now do not happen on Zoom. So, I also suspect the “so what” here is limited.”

The paragraph above makes several claims, and we address each in turn:


1. Low correlation between SES and access to high-speed internet access

In contrast to your intuition, there is ample evidence that lower SES is associated with worse internet quality. For example, researchers found that 92% of the large US cities they examined had disparities based on income¹³. Another report interviewed residents of poor, non-white US neighborhoods wherein they described internet speeds “so slow that [they don’t] meet Zoom’s recommended minimum for group video calls.” This same report noted that a wealthier neighborhood in New Orleans had internet speeds 400x faster than a middle-class neighborhood¹². We can also consider patterns outside urban areas. As you mention, rural areas—which often experience widespread poverty—are also subject to poor internet access and connectivity^{9,32}.

Broadly, low SES groups and rural communities are only two of the population groups that are disadvantaged in terms of internet access. Internet quality is correlated with a variety of demographic and psychographic variables, including age and race among others⁹. If these segments systematically suffer from poorer internet speeds, they are therefore at higher risks for glitches and the disadvantages that glitches create. Taken together, we believe the policy implications of our work are wide-reaching and apply to diverse segments of the population. We have bolstered these arguments beginning on page 21 of the revised manuscript.

2. By 2024 most low income individuals are working and living fully in person, with little zoom activity.

Although remote work is common among wealthier employees, lower-income individuals use videoconference for a variety of other common practices outside of just the professional. As our Extended Data (Figure 1) illustrates, teleconferencing is prevalent across a wide range of other important life domains, each of which shape the everyday lives of not just wealthy remote workers, but individuals across life stages, ages, economic statuses, and needs. Below is a snapshot of the figure:



What's more, videoconferencing has been promoted as an explicit solution to problems that low-income (and other underserved communities) experience. For example, videoconferencing is crucial for low-income Americans to access healthcare, and extensive efforts are being made to connect telehealth providers and underserved communities^{33,34}. (Notably, prior to COVID, telehealth was primarily targeting those "located in a rural area" or where there was a "health provider shortage"³⁵.) Furthermore, regardless of whether job-seekers are applying for an in-person or remote/hybrid job, the *majority* of Americans begin the recruitment process (i.e., early-round interviews) via videoconferencing⁸.

3. Court proceedings now do not happen on Zoom

Related to the above point, there is substantial evidence of videoconferencing adoption in many legal proceedings even several years post-pandemic⁴⁻⁶. For example, in Washington State, approximately 1 in 3 civil and criminal cases are handled via videoconference as of 2023³⁶. Many other states across the country have similarly adopted videoconferencing in their court hearings, ranging from criminal and civil trials, to traffic violations, workers' compensation, divorce and mediation proceeding, and child support hearings. Even though we no longer need to socially distance for public health reasons, this shift to virtual hearings likely persists because it offers multiple benefits relating to equitable access (e.g., minimizing transportation costs, enhancing public transparency)³.

The chart we reference above provides additional helpful statistics, and the revised manuscript reflects our arguments.

"3) Correlation vs Causation. I was concerned with omitted variable bias for studies 1 and 5. The types of people that might report "glitches" or say "freeze" in zoom calls may be more liable to complain in general. Relaxed positive individuals likely complain less.

So, their calls are not going to get tagged as “glitchy”. More critical negative individuals probably complain more about glitches so get tagged. But these same negative critical people likely give more critical call evaluations. Hence, it is not surprising that people that complain a lot about call quality also complain a lot about call content, generating a correlation between “glitches” and bad call evaluations. Studies 2 to 4 are more convincing as RCTs, although the magnitude of the induced glitches seem very large for such a short call so it’s more than brief glitches being examined here.”

We included Studies 1 and 6 because we believe that real-world, ecologically valid illustrations are valuable. We opted to leverage existing (albeit correlational) datasets to document meaningful patterns, rather than pursue an RCT and attempt to induce negative causal effects for real-life parolees. But of course, given their observational nature, these studies are limited in their ability to shed light on causality. For this reason, we offer the broader body of evidence, which currently includes these 2 archival datasets and 6 carefully controlled experiments.

To the question of glitch magnitude, this does not appear to have a straightforward relationship with negative downstream effects. For example, in our new Study 4, even though extended freezes were twice as long as fleeting freezes and were perceived as more disruptive, they were significantly less uncanny and had a smaller negative effect on interpersonal judgment.

At the same time, one could argue that many of our glitchy stimuli are *conservative* relative to real-world glitches. In many of our studies, we carefully timed glitches to occur during natural speech pauses, ensuring they didn’t interrupt the actual conversation content or prevent calls from proceeding to their natural conclusion. These conditions are not necessarily afforded in many real-world video calls.

You also raise a specific concern about negative individuals—that they may be more likely to report or call out glitches and this explains the observed correlations. While plausible, in Study 1, we find that glitches are associated with lower evaluations even after controlling for participants’ pre-conversation affect (a proxy for an individual’s predisposition for negativity). This suggests that the observed relationship between glitches and communication outcomes is not limited to such negative individuals. Similarly, in Study 6, our effects are not impacted by variance stemming from primary or secondary parole commissioner, also suggesting that individual dispositions are not central.

“4) Inconsistencies: One final thing is the text and the methods section are at times inconsistent. For example, Study 2 was described in the text as having “four brief (<2 second)” freezes, but in the methods as “five freezes...from 0.5 to 2 seconds in duration”. This seems like a material difference given this is a 75 second video. I would like to know the precise length of all the freezes and the total length of the call, ideally we examples to show this. Also, some of the studies lacked information on how long the freezes lasted for (study 4 for example).”

Thank you for raising the issue of reporting inconsistencies. We have taken care to ensure that all reporting of glitches is correct. Further, for additional ease of processing, below is a table delineating the glitches used in each study (timestamp and duration) and a link to the video

stimuli. Note that this table is meant only for internal review purposes, as the human actors have not offered their permission to publish their likenesses.

We hope the myriad changes we have made throughout the paper have clarified the surprisingness and robustness of our findings and the breadth and importance of the implications. Thank you for your time and suggestions.

[REDACTION]

This was a table detailing the glitches in the main study that was redacted because it includes YouTube links to stimuli that we do not have the rights to publish.

References:

- 1 Iowa Board of Parole. Attending a parole board hearing. <https://bop.iowa.gov/hearings/attending-parole-board-hearing> (2024).
- 2 State of Connecticut. State of connecticut board of pardons and paroles. <https://portal.ct.gov/bopp> (2024).
- 3 Louisiana Department of Public Safety & Corrections. Louisiana board of pardons and committee on parole. <https://doc.louisiana.gov/imprisoned-person-programs-resources/pardons-parole/> (2024).
- 4 Reed, A. Virtual court hearings earn permanent spot after pandemic's end. <https://news.bloomberglaw.com/us-law-week/virtual-court-hearings-earn-permanent-spot-after-pandemics-end> (2023).
- 5 Thomsen, J. Virtual court hearings are here to stay post-pandemic, survey finds. *Natl. Law J.* <https://www.law.com/nationallawjournal/2021/08/18/virtual-court-hearings-are-here-to-stay-post-pandemic-survey-finds> (2021).
- 6 Pattison-Gordon, J. For some states, remote court hearings are here to stay. <https://www.govtech.com/public-safety/for-some-states-remote-court-hearings-are-here-to-stay> (2023).
- 7 Definitive Healthcare. Hospital telehealth adoption by state. <https://www.definitivehc.com/resources/healthcare-insights/hospital-telehealth-adoption-by-state> (2024).
- 8 Hankel, I. The virtual interview is the new resume: What you need to know. *Forbes*. <https://www.forbes.com/sites/forbesbusinesscouncil/2022/05/05/the-virtual-interview-is-the-new-resume-what-you-need-to-know/?sh=3e3acd4f486d> (2022).
- 9 Pew Research Center. Internet, broadband fact sheet. <https://www.pewresearch.org/internet/fact-sheet/internet-broadband/?menuItem=89fe9877-d6d0-42c5-bca0-8e6034e300aa> (2024).
- 10 Carnegie Corporation of New York. New data shows how digital inequality affects lower-income students and their families. <https://www.carnegie.org/our-work/article/new-data-shows-how-digital-inequality-affects-lower-income-students-and-their-families/> (2021).
- 11 Insight Policy Research. ConnectHome initiative: Final report. <https://www.huduser.gov/portal/publications/ConnectHome.html> (2018).
- 12 Yin, L. & Sankin, A. Investigation finds lower internet speeds for higher prices in poor, less white U.S. Neighborhoods. <https://www.pbs.org/newshour/economy/investigation-finds-lower-internet-speeds-for-higher-prices-in-poor-less-white-u-s-neighborhoods> (2022).
- 13 Mohammad, L., Kenin, J. & Chang, A. Your internet is slow because of where you live, not what plan you buy. <https://www.delawarepublic.org/2022-10-26/your-internet-is-slow-because-of-where-you-live-not-what-plan-you-buy> (2022).
- 14 Howell, C. & West, D. M. The internet as a human right. <https://www.brookings.edu/articles/the-internet-as-a-human-right/> (2016).
- 15 Human Rights Council. The promotion, protection and enjoyment of human rights on the internet. https://www.article19.org/data/files/Internet_Statement_Adopted.pdf (2016).
- 16 Short, J., Williams, E. & Christie, B. *The Social Psychology of Telecommunications*. (Wiley, 1976).

- 17 Daft, R. L. & Lengel, R. H. Organizational information requirements, media richness and structural design. *Manag. Sci.* **32**, 554-571 (1986).
- 18 Kruger, J., Epley, N., Parker, J. & Ng, Z.-W. Egocentrism over e-mail: Can we communicate as well as we think? *J. Pers. Soc. Psychol.* **89**, 925-936 (2005).
- 19 Gray, K. & Wegner, D. M. Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition* **125**, 125-130 (2012).
- 20 Wang, S., Lilienfeld, S. O. & Rochat, P. The uncanny valley: Existence and explanations. *Rev. Gen. Psychol.* **19**, 393-407 (2015).
- 21 Song, S. W. & Shin, M. Uncanny valley effects on chatbot trust, purchase intention, and adoption intention in the context of e-commerce: The moderating role of avatar familiarity. *Int. J. Hum. Comput. Interact.* **40**, 441-456 (2024).
- 22 Schroeder, J., Kardas, M. & Epley, N. The humanizing voice: Speech reveals, and text conceals, a more thoughtful mind in the midst of disagreement. *Psychol. Sci.* **28**, 1745-1762 (2017).
- 23 Dennis, A. R., Fuller, R. M. & Valacich, J. S. Media, tasks, and communication processes: A theory of media synchronicity. *MIS Q.* **32**, 575-600 (2008).
- 24 Lombard, M. & Ditton, T. At the heart of it all: The concept of presence. *J. Comput.-Mediat. Commun.* **3** (1997).
- 25 Bagozzi, R. P., Yi, Y. & Phillips, L. W. Assessing construct validity in organizational research. *Admin. Sci. Q.* **36**, 421-458 (1991).
- 26 Schein, C. & Gray, K. The eyes are the window to the uncanny valley: Mind perception, autism and missing souls. *Interact. Stud.* **16**, 173-179 (2015).
- 27 Cohen, J. *Statistical Power Analysis for the Behavioral Sciences* (Erlbaum, 1969).
- 28 Eaton, K. Zoom's vision pro app shows your company's strange, 3-D virtual meeting future. *Inc.* <https://www.inc.com/kit-eaton/zooms-vision-pro-app-shows-your-companys-strange-3d-virtual-meeting-future.html> (2024).
- 29 Tinwell, A., Grimshaw, M. & Williams, A. The uncanny wall. *Int. J. Arts Technol.* **4**, 326-341 (2011).
- 30 Ratajczyk, D. Shape of the uncanny valley and emotional attitudes toward robots assessed by an analysis of youtube comments. *Int. J. Soc. Robot.* **14**, 1787-1803 (2022).
- 31 Simonsohn, U. Groundhog 2.0: Further addressing the threat R poses to reproducible research. <https://datacolada.org/100> (2022).
- 32 Vogels, E. Digital divide persists even as americans with lower incomes make gains in tech adoption. <https://www.pewresearch.org/fact-tank/2021/06/22/digital-divide-persists-even-as-americans-with-lower-incomes-make-gains-in-tech-adoption/> (2021).
- 33 Adetosoye, F., Hinton, D., Shenai, G. & Thayumanavan, E. Virtual health for all: Closing the digital divide to expand access. <https://www.mckinsey.com/industries/public-sector/our-insights/virtual-health-for-all-closing-the-digital-divide-to-expand-access> (2023).
- 34 Koehn, A. Apps and technology help low-income people access healthcare. <https://news.medill.northwestern.edu/chicago/apps-and-technology-help-low-income-people-access-healthcare/> (2016).
- 35 American Hospital Association. Fact sheet: Telehealth waivers. <https://www.aha.org/fact-sheets/2025-02-07-fact-sheet-telehealth-waivers> (2025).
- 36 Larsen, P. Summary report of remote proceedings prevalence and practices in Washington courts. https://www.courts.wa.gov/programs_orgs/pos_bja/RemoteTF (2023).