NBER WORKING PAPER SERIES

HOW WILL GENERATIVE AI IMPACT COMMUNICATION?

Joshua S. Gans

Working Paper 32690
http://www.nber.org/papers/w32690

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024

Joshua Gans has drawn on the findings of his research for both compensated speaking engagements and consulting engagements. He has written the books Prediction Machines, Power & Prediction, and Innovation + Equality on the economics of AI for which he receives royalties. He is also chief economist of the Creative Destruction Lab, a University of Toronto-based program that helps seed stage companies, from which he receives compensation. He conducts consulting on anti-trust and intellectual property matters. He also has equity and advisory relationships with a number of startup firms. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

How will Generative AI impact Communication?
Joshua S. Gans
NBER Working Paper No. 32690
July 2024
JEL No. D83,O31

ABSTRACT

This paper examines the impact of Generative AI (GAI) on communication through the lens of salience and signalling models. It explores how GAI affects both senders' ability to create salient messages and receivers' costs of absorbing them. The analysis reveals that while GAI can increase communication by reducing costs, it may also disrupt traditional signalling mechanisms. In a salience model, GAI generally improves outcomes but can potentially reduce receiver welfare. In a pure signalling model, GAI may hinder effective communication by making it harder to distinguish high-quality messages. This suggests that GAI's introduction necessitates new instruments and mechanisms to facilitate effective communication and quality assessment in this evolving landscape.

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
and NBER
joshua.gans@rotman.utoronto.ca

1 Introduction

Generative artificial intelligence (GAI) has provided a set of tools that dramatically lower the costs of composing text (e.g., OpenAI's ChatGPT or Anthropic's Claude), creating images (e.g., Dall-E, Google Gemini or Midjourney), and potentially more. Each of these is part of communication activities in which a sender wants to send a message to a receiver (or receivers) that they aim for the reader to absorb, learn from, or appreciate. At the same time, GAI tools, by making communication easier, also assist receivers in potentially reducing the costs of absorbing messages. This might be achieved because the messages are clearer and more persuasive, but also because GAI allows people with poor skills in the receiver's language to generate messages that are easier to parse.

Such communication faces a fundamental challenge in that receivers typically have scarce attention or other costs from absorbing messages; that is, one might send a receiver a message, but it takes additional costly effort for the receiver to actually absorb that message.[1] This can be an issue even if the sender and receiver have congruent interests in having the receiver absorb the message.[2] If the sender can signal the importance of a message, then this can provide an incentive to a receiver to read a particular message.

This challenge was examined in an advertising context by which consumers (the receivers here) could observe a signal of the cost expended in creating an advertisement, and this allowed advertisers with higher-quality products (the senders) to generate a more salient signal (Milgrom and Roberts, 1986). There is, of course, a conflict of interest between the sender and receiver in advertising that creates a specific signalling issue. At the same time, spending more on advertising can also, in some cases, lead to ads that are more attention-grabbing or salient to consumers, which can be seen as reducing the costs to consumers of absorbing ads. As already noted here, the impact of GAI impacts the attention costs of communication on both the *signalling* and the *salience* margins.

2 A Simple Salience Model

To explore these effects, consider a situation where a sender (S) wants to send a message to a receiver (R). The sender's message can take values of $\{1, \varnothing\}$ corresponding to a sent message and no message, respectively. When presented with a message, the receiver can choose $a_R \in \{0, 1\}$ where $a_R = 1$ $(= 0)$, which involves the R absorbing (ignoring) the

[1] Cremer et al. (2007) and Garicano and Prat (2013) provide analyses of such costly communication within organisations. See also Bilancini and Boncinelli (2018).

[2] Dewatripont and Tirole (2005) provide a model where there are such costs of sending and receiving messages and where congruency of interests impacts the total amount of information communicated.

message. The receiver faces a cost c_R of absorbing the message and receives value v_R from the message's content. The sender faces a cost c_S of sending the message and obtains value v_S if the receiver absorbs the message.

v_R is unknown to the receiver prior to absorbing a message. Let p be the probability that $v_R = v_R^h$ and suppose that $v_R^h > c_R > v_R^l = 0$. Therefore, absent a signal, R will absorb a message if and only if $pv_R^h \geq c_R$.[3] If this condition holds, all messages get through, and there is no issue on the sender side. What if, however, $pv_R^h < c_R$? In this case, no messages are absorbed.

The sender's value if a message is absorbed is a_S which can take on values of v_S^h and v_S^l depending on the value of the message to the receiver. The sender knows this value prior to sending the message. In this initial model, is is assumed that a sender can choose to invest $a_S \in \{I, 0\}$, where $a_S = I$ involves a cost of c_S but allows the cost to the receiver of absorbing the sender's message to be reduced by $\Delta_R > 0$. This salience approach is contrasted with a (signalling) model presented in Section 3 below where, at a cost of e, the sender can observably expend effort in sending a message.

2.1 Outcome without GAI

With respect to investment in improving salience, there exist equilibrium outcomes involving no communication, partial communication, and full communication (with and without investments in salience). The following proposition characterises these equilibrium outcomes.

Proposition 1 *Let $\{a_S^h, a_S^l\}$ be the pure strategy (perfect Bayesian) equilibrium investments of the h- and l-sender types, respectively. In each equilibrium, messages are sent by senders regardless of type.*

1. *(No communication) If (i) $pv_R^h < c_R$ and $v_S^h < c_S$ or (ii) $v_R^h < c_R - \Delta_R$, then $\{a_S^h, a_S^l\} = \{0, 0\}$ and R does not absorb any message sent;*

2. *(Partial communication) If (i) $pv_R^h < c_R$; (ii) $v_S^h \geq c_S$; (iii) $v_R^h \geq c_R - \Delta_R$; and (iv) $v_S^l < c_S$, then $\{a_S^h, a_S^l\} = \{I, 0\}$ and R absorbs salient messages only;*

3. *(Full communication, investment) If (i) $pv_R^h \in [c_R - \Delta_R, c_R)$ and (ii) $v_S^l \geq c_S$, then $\{a_S^h, a_S^l\} = \{I, I\}$ and R absorbs salient messages only; and*

4. *(Full communication, no investment) If $pv_R^h \geq c_R$, then $\{a_S^h, a_S^l\} = \{0, 0\}$ and R absorbs all messages.*

[3] It is assumed here that when R is indifferent between $a_R = 1$ and $a_R = 0$, they choose $a_R = 1$.

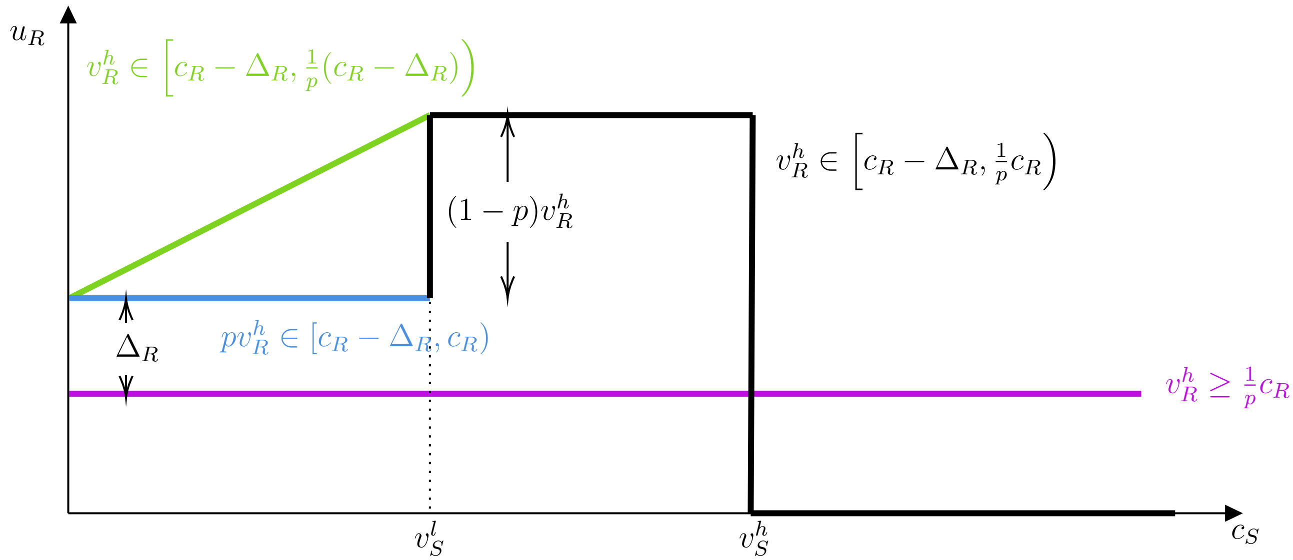

Figure 1: Expected Utility of the Receiver in Equilibrium

The conditions for the no communication and full communication equilibrium outcomes are quite intuitive. No communication occurs if the h-senders do not find it worthwhile to invest in salience and R does not find it worthwhile to absorb the average message. In the full communication equilibrium outcomes, either R finds it optimal to absorb messages of both types that are not salient or if they are salient, both types find it optimal to invest in salience.

The interesting outcome here is where R only absorbs salient messages, and only the h-senders find it optimal to send such messages. In this case, R absorbs salient messages because it knows that messages are only valuable when salient. In other words, not only are salient messages easier to absorb, but there is also signalling value in the fact that it is costly for senders to invest in salience. Note that this equilibrium outcome results in the highest expected utility for the receiver; see Figure 1.

However, note that there is also a mixed-strategy equilibrium. That arises under conditions (i) - (iii) of the partial communication outcome but when $v_S^l \geq c_S$. In this case, if R only absorbs salient messages, either $pv_R^h \geq c_R - \Delta_R$ and the full communication equilibrium with investment exists or $pv_R^h < c_R - \Delta_R$ and the full communication equilibrium does not exist. In this case, the pure strategy partial equilibrium does not exist because the l-senders would find it optimal to send messages if the receiver believed that all salient messages were of high quality. The only equilibrium that exists is where l-receivers play a mixed strategy, with some fraction of them investing in salience and the others not. This outcome is depicted in Figure 1 for the case where $v_R^h \in [c_R - \Delta_R, \frac{1}{p}(c_R - \Delta_R))$. l-senders mix just enough that R's posterior probability that a salient message is of high quality

2.2 Outcome with GAI

What is the impact of GAI on these equilibrium outcomes? GAI makes it easier for senders to generate messages that are more salient. That is, GAI results in a reduction in c_S (or equivalently an increase in Δ_R). A reduction in c_S makes it less likely that a no-communication equilibrium exists by making it easier for h-senders to invest in salience. This makes a partial communication more likely to arise when c_S is initially high but there is a mitigating effect in that GAI also makes it easier for l-senders to invest when $c_S \leq v_S^h$ initially. As can be seen in Figure 1, as c_S falls, receiver utility rises, but then with further reductions, it can fall. By contrast, c_S reductions are always good for senders. Nonetheless, GAI always increases the total amount of communication.

This approach focuses on GAI that might assist in reducing the costs to receivers of absorbing information, but it does so in a context where those costs are somewhat protective of receivers. Senders could invest in reducing those costs, which would increase the chances of their messages being communicated to receivers, but, at the same time, that acted as a screening device by selecting for messages that senders placed a relatively high value on being received. A key assumption in the model here was that there was some alignment of interests in that receivers only valued the messages from high-valued senders. If, alternatively, the assumptions were reversed – and senders with messages of low receiver value also had a relatively high value of their messages being received – then salience may send an 'anti-signal' to receivers. In this situation, there would be no partial communication outcome, with receivers choosing to accept or reject all messages. That said, GAI, by reducing receiver costs, would still encourage more communication.

It is also possible that GAI could have an impact on p, the probability that a message was of high value to the receiver. This could be achieved by personalising messages or other ways of making messages more useful to the receiver. It is easy to see that an increase in p would both promote more communication and also improve receiver welfare (locally) as the messages became more relevant to them. That said, an increase in p could also reduce the need for senders to make investments to improve salience, which may change the type of equilibrium outcome in a way that leaves receivers worse off.

3 A Pure Signalling Model

Given the analysis thusfar, it is useful to consider an alternative approach that focuses on the ability of the sender to signal the quality of their message to the receiver more directly. The previous model embedded a signal within a mechanism that made messages easier for receivers to absorb. However, imagine if the sender could send a signal that was indicative of the amount of personal effort or "work" they put into the message itself. To this end, suppose that a sender can choose their effort, e, at a cost of e to generate the message. With probability $\rho(e)$, an increasing and concave function in e where $\rho'(0) > 0$, the receiver learns the e associated with a message, while with probability $1 - \rho(e)$, the receiver is unsure of the effort level. In what follows, we will model GAI as making it easier to generate a message that looks like it has taken effort, thereby jamming what might otherwise be a quality signal.

3.1 Outcome without GAI

Without GAI, there exist two pooling equilibrium outcomes with full or no communication, respectively, but also a separating equilibrium outcome involving partial communication. The following proposition characterises these.

Proposition 2 *Let $\{e_S^h, e_S^1\}$ be the (perfect Bayesian) equilibrium efforts of the h- and l-sender types, respectively. In each equilibrium, messages are sent by senders regardless of type.*

1. *(No communication) If $v_R^h < c_R$, then $\{e_S^h, e_S^l\} = \{0, 0\}$ and R does not absorb any message sent;*

2. *(Partial communication) If (i) $pv_R^h < c_R$; (ii) $v_R^h \geq c_R$; and (iii) $\rho(\hat{e})v_S^h \geq \hat{e}$, then $\{e_S^h, e_S^1\} = \{\hat{e}, 0\}$ where $\rho(\hat{e})v_S^l \approx \hat{e}$, are the equilibrium effort levels and R absorbs any message that receives an 'effort' signal of at least $\hat{e}$; and*

3. *(Full communication) If $pv_R^h \geq c_R$, then $\{e_S^h, e_S^l\} = \{0, 0\}$ and R absorbs all messages.*

The interesting equilibrium outcome is the separating equilibrium. In that equilibrium, the h-sender chooses an effort level, $\hat{e}$ that is so low enough that the l-sender type does not match the h-sender's effort level of $\hat{e}$; that is, $\rho(\hat{e})v_S^l < \hat{e}$. At this effort level, an h-sender's incentive compatibility constraint does not bind. The l-sender cannot get away with choosing a lower effort level as the receiver will not absorb messages below a threshold of $\hat{e}$ even if the effort is revealed to them. Moreover, because $\rho(e)$ is increasing, a message without an effort signal has a posterior probability of being a high-valued message that is lower than the prior probability, and so the receiver will not absorb such messages. This implies that a separating equilibrium exists if $v_R^h \in [c_R, \frac{1}{p}c_R\}$.

3.2 Outcome with GAI

The impact of GAI is to make it harder for any sender to signal their effort level; that is, $\rho(e)$ is lower as is $\rho'(e)$. This will only impact the separating equilibrium with partial communication. This means that the level of $\hat{e}$ that allows the l-sender's incentive compatibility constraint to bind will be higher. This reduces the h-sender's payoff. For the receiver, their welfare is determined by whether GAI leads to an increase or decrease in the probability of receiving an 'effort' signal. The GAI itself reduces that probability but, at the same time, changes the amount of effort h-senders expend, which can mitigate that effect.

To explore this, suppose that the probability that a receiver observes effort is $\mu\rho(e)$ where GAI reduces μ. We are interested in the impact of μ on $\mu\rho(\hat{e}(\mu))$. As $\mu\rho(\hat{e}(\mu))v_S^l = \hat{e}$,

$$\frac{d\mu\rho(\hat{e}(\mu))}{d\mu} = \rho(\hat{e}(\mu)) + \frac{\partial\rho}{\partial e}\frac{d\hat{e}}{d\mu} > 0$$

as $\frac{d\hat{e}}{d\mu} = -\frac{\rho(e)}{\mu\rho'(e)v-1} > 0$ if $\mu\rho'(e)v < 1$. Under this condition, a reduction in μ leads to a reduction in communication in equilibrium.[4] Thus, the amount of communication absorbed by the receiver falls.

In summary, in contrast to the salience model, where receivers might gain from GAI and there is an increase in communication, the pure signalling model does not facilitate communication and is does not help senders or receivers. This is because GAI removes the ability for actions that allow receivers to more clearly judge the quality of incoming messages.

4 Conclusion

All of this points to the need for new instruments to assist in communication when GAI is available. When signalling, in particular, is important, GAI will increase demand for richer signals and also other mechanisms (including reputational mechanisms) to assist receivers in judging the quality of messages they might absorb.

Recent empirical evidence by Cowgill et al. (2024) provides support for the theoretical framework presented in this paper. Their study on the impact of generative AI (specifically ChatGPT) in entrepreneurship and hiring contexts reveals that AI can indeed reduce the informational value of signals, aligning with the signalling model. They found that ChatGPT introduces about a 2 percent information loss by increasing noise in evaluations, and both senders and receivers respond by demanding more costly signals. Interestingly, this effect

[4]This can be confirmed by a simple example where $\rho(e) = \mu\sqrt{e}$ leads to $\mu\sqrt{\hat{e}}v_S^l = \hat{e} \implies \mu\sqrt{\hat{e}} = \mu^2 v_S^l$ implying that a reduction in μ lowers the equilibrium $\rho(\hat{e})$.

is more pronounced for senders from lower-income countries and those with lower signalling ability, suggesting that GAI acts as a substitute for signalling ability. The study also highlights the importance of cultural context, as receivers from cultures with high uncertainty avoidance and long-term orientation showed greater sensitivity to ChatGPT use. While their research did not find significant evidence of non-experts successfully faking expertise using ChatGPT, it underscores the complex ways in which GAI is reshaping communication and signalling in economic transactions. These findings reinforce an implication of the present paper for the need to develop new instruments and mechanisms to facilitate effective communication and quality assessment in the era of generative AI.

References

Bilancini, E. and Boncinelli, L. (2018). Signaling with costly acquisition of signals. *Journal of Economic Behavior & Organization*, 145:141–150.

Cowgill, B., Hernández-Lagos, P., and Wright, N. L. (2024). Does ai cheapen talk? evidence from global startup and hiring contexts. *Working Paper Available at SSRN 4702114*.

Cremer, J., Garicano, L., and Prat, A. (2007). Language and the theory of the firm. *The Quarterly Journal of Economics*, 122(1):373–407.

Dewatripont, M. and Tirole, J. (2005). Modes of communication. *Journal of Political Economy*, 113(6):1217–1238.

Garicano, L. and Prat, A. (2013). Organizational economics with cognitive costs. *Advances in Economics and Econometrics*, 1:342.

Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821.