

Will ChatGPT Break Notice and Comment for Regulations?

In brief...

ChatGPT, an artificial intelligence (AI) tool based on a large language model, has the potential to affect the notice-and-comment process for rulemaking. This commentary considers how ChatGPT might be used to generate public comments and explains the existing safeguards for submitting public comments that prevent the system from being overwhelmed.

By: Mark Febrizio | January 13, 2023

A new AI tool is making waves among researchers, educators, and internet commentators, with some being excited about its potential while others are concerned about detrimental effects. ChatGPT, a project of [OpenAI](#), is a chatbot that relies on a large language model and machine learning techniques to produce answers to user-submitted queries. Since its launch in late November 2022, the chatbot has dominated internet conversations as interested individuals from various backgrounds have [experimented](#) with generating stories, reviewing software code, and summarizing complex topics.

Unsurprisingly, researchers are most intrigued with how ChatGPT might affect their own domains; public policy scholars are no different. One article [theorized](#) about how ChatGPT could break the public comment process if it were used to “flood the system” with computer-generated submissions. Certainly, there is much to learn about ChatGPT and comparable AI tools. But given the safeguards and practices related to commenting on federal agency rules, ChatGPT seems unlikely to introduce novel problems to the notice-and-comment process.

An Experiment with ChatGPT

A first step of investigating the likelihood of encountering public commenting concerns, is ChatGPT capable of mimicking public submissions to agency rulemakings? Said differently, can the AI tool effectively create public comments similar to those submitted by members of the public? In short, yes; it can.

I asked ChatGPT to write a public comment within 250 words that objects to a [proposed rule](#) from the Department of Labor on classifying independent contractors. With minimal prompting, the chatbot produced a convincing submission from a member of the public:

As a concerned member of the public, I strongly object to the proposed rule titled "Employee or Independent Contractor Classification under the Fair Labor Standards Act" that can be accessed at the link provided.

This proposed rule would make it easier for companies to classify their workers as independent contractors instead of employees, depriving these workers of essential protections and benefits that are afforded to employees under the Fair Labor Standards Act. This includes minimum wage, overtime pay, and protection from discrimination.

Furthermore, classifying workers as independent contractors instead of employees allows companies to avoid paying their fair share of taxes, which undermines the social safety net and puts an unfair burden on taxpayers.

Additionally, this proposed rule would disproportionately impact vulnerable workers who are more likely to be classified as independent contractors, such as those in low-wage industries and gig workers. This would only further exacerbate income inequality and weaken the labor market.

Overall, this proposed rule is misguided and harmful to workers, taxpayers, and the economy as a whole. I urge the Department of Labor to withdraw this proposal and instead focus on strengthening protections for workers and promoting fair labor practices.

Further, ChatGPT was able to easily pivot to the opposite perspective:

As a concerned member of the public, I strongly support the proposed rule titled "Employee or Independent Contractor Classification under the Fair Labor Standards Act."

This proposed rule would clarify the criteria for determining whether a worker should be classified as an employee or an independent contractor under the Fair Labor Standards Act. This clarity is essential for ensuring that workers are properly classified and receive the appropriate protections and benefits.

Furthermore, this proposed rule would provide greater flexibility for workers who prefer to be classified as independent contractors. This would allow them to have more control over their schedules and work arrangements, which can be beneficial for those who value the autonomy and freedom that comes with being an independent contractor.

Additionally, this proposed rule would provide greater certainty for businesses that are unsure of how to properly classify their workers. This would reduce the potential for costly litigation and allow businesses to focus on their operations and growth.

Overall, this proposed rule is well-reasoned and would benefit both workers and businesses. I urge the Department of Labor to move forward with this proposal and provide much-needed clarity and flexibility in the classification of workers under the Fair Labor Standards Act.

As someone who has read through hundreds of public comments as part of a [previous study](#), these responses stand out to me as comparable to mass comment campaign (MCC) submissions—a common type of public comment on agency proposals. Scholars define [MCCs](#) as “identical and near-duplicate comments sponsored by organizations and submitted by group members and supporters to government agencies in response to proposed rules.” Further, it makes sense that MCCs are a good testing ground for ChatGPT’s abilities because they are numerous in volume, relatively short, and often pointed in tone.

Implications of ChatGPT's Comment Generation Abilities

Evidently, ChatGPT is capable of producing comments that resemble those submitted by concerned individuals. However, is that a problem?

Importantly, using a tool to help compose the content of public submissions, whether sponsored by organizations or not, is not inherently a problem for notice and comment. Individuals and organizations are free to use various tools and resources to create comments that represent their views. From the agency perspective, this is not a big concern because the government’s [obligation](#) is to identify and consider significant issues and offer a reasoned explanation of how they influenced its decision. In addition, some agencies, [such as the Environmental Protection Agency](#) (EPA), are already accustomed to [identifying](#) large sets of duplicate or substantially similar comments and responding to them in turn. While concerns about MCCs are often linked to computer-generated or “malattributed” public submissions, scholars have [concluded](#) that they “do not represent a crisis for the regulatory state at this time.” In a sense, the potential issues posed by AI chatbots are not novel.

But an open question, arguably, is whether ChatGPT can aid in creating a set of unique or sufficiently distinctive comments that elude agency tools that identify MCCs. When I asked ChatGPT about its theoretical potential to generate unique comments, it responded affirmatively:

As a language model AI, I have the capability to generate an almost limitless number of unique and distinct comments on the proposed rule. However, this would depend on the specific rule and the context surrounding it. The more information that is available on a topic, the more I can generate variations of comments based on the provided information.

But even assuming that ChatGPT has an accurate perception of its own abilities, are existing protections sufficient to mitigate threats to the functioning of notice and comment?

Safeguards In Public Commenting

To better investigate the issue at hand, assume that ChatGPT facilitates a virtually unlimited supply of AI-generated public comments for any proposed rule. Would that threaten the functioning of notice and comment for federal regulations? I would argue that the safeguards in place—at least on [Regulations.gov](#), the primary site for receiving public submissions—are sufficient to prevent AI-generated content from overwhelming the system.

The first key safeguard is that commenters must go through a [Completely Automated Public Turing test to tell Computers and Humans Apart](#) (CAPTCHA) system when submitting comments through the web [Regulatory Studies Center](#)

user interface. Most regular users of the internet will be familiar with CAPTCHA tests (e.g., a checkbox saying “I’m not a robot” paired with identifying which pictures in a 3x3 grid contain a crosswalk), which seek to distinguish human users from automated ones. Specifically, Regulations.gov [relies on](#) Google’s [reCAPTCHA](#) “to support the integrity of the rulemaking process and manage the role of software-generated comments.”

Second, commenters can also make submissions through the [Regulations.gov Application Programming Interface](#) (API), which is managed by the General Services Administration (GSA). In particular, using the POST API makes sense for organizations or sponsors that have collected and want to submit large batches of comments. Similar to the web user interface, the API was designed with its own safeguards.

To gain access to the POST API, organizations must register for API keys (they can have [up to two](#)). But obtaining an API key does not permit unlimited access. Usage is restricted by rate limits for each key, with the commenting API having a [rate limit](#) of “50 requests per minute with a secondary limit of 500 requests per hour.” This means that an API key making 50 submissions per minute would be constrained by the hourly rate limit after 10 minutes. Also, a user trying to submit 500 comments within one or two minutes would be constrained by the primary rate limit. Finally, GSA can [revoke](#) a user’s access if malicious activity is detected.

Even with an unlimited supply of AI-generated content, a malicious user would quickly hit a bottleneck when trying to submit those comments on agency rules. The web user interface was not designed for submitting large batches of comments, and the reCAPTCHA system is built to preclude computer-based tools from accessing and making numerous submissions in an automated manner. The commenting API requires pre-authorization and limits how many submissions a key holder can make in a certain amount of time. If GSA detects questionable activity, it can identify which key holders are responsible and revoke access. The upshot is that public commenting appears to be fairly resilient to technology-related abuses of the system.

Technology has Both Benefits and Costs

In response to these arguments, one might object that a sufficiently smart or motivated person or group could get around these safeguards. And they are likely correct. Further, although Regulations.gov is the primary platform used in the federal notice-and-comment process, it is not the only one. Agencies that rely on their own public commenting portals (or state-level rulemaking processes) could be more vulnerable to an influx of AI-generated comments. Here, it would be prudent for GSA to consider how the expansion of chatbot technologies might affect the notice-and-comment process and the operations of Regulations.gov. This [2021 final report](#) on Mass, Computer-Generated, and Fraudulent Comments, which was prepared for the Administrative Conference of the United States (ACUS) by leading scholars, is a good place to begin.

The reality is that we can never remove all risk from a system—especially when humans are involved. New technologies may alter the risk management considerations, but they are often not the fundamental

source of the problem. Technology is rarely inherently bad or good, rather its impact depends on whether it is used in beneficial or malicious ways.

In other words, new tools imply tradeoffs—a corollary to the benefit-cost framework used in regulatory impact analysis. ChatGPT may produce effects that cut both ways. While AI-generated content may pose risks, agencies can also benefit from advances in machine learning techniques to identify MCCs and computer-generated comments. One [emerging example](#) is [GPTZero](#), which uses its own model to differentiate between human- and AI-generated content.

As a quick test, I input the two comments described earlier that were created by ChatGPT, and GPTZero correctly spit out: “Your text is most likely to be AI generated!” Because every experiment (even cursory ones) deserves a control group, I also fed the opening paragraphs of my commentary explaining the [revamped Regulations.gov](#) platform to GPTZero. The result: “Your text is likely human generated!” Pretty cool.

While ignoring the problems that come with new technologies would be a mistake, having a reasoned response that takes account of existing safeguards and considers both potential benefits and costs is also important. Based on what we know, public commenting is unlikely to break because of ChatGPT.