

Discussion of Blog Post: The Continuing Vitality of Epicurean Physics

Post by “Cassius” of April 27, 2026 at 9:54 AM

This thread is for discussion of the blog article:



[The Continuing Vitality Of Epicurean Physics](#)

Introduction: A Story About the Sun That Is Really About Everything

In a recent scholarly essay, T.H.M. Gellar-Goad takes up one of the most ridiculed positions in the history of ancient philosophy: the Epicurean claim that the sun is the size it appears to be. Cicero mocked it. Stoic philosophers used it as evidence that Epicureans were intellectually unserious. Modern critics have repeated the mockery with updated vocabulary. And on the surface, the laughter seems justified — after all, we...



Cassius

April 27, 2026 at 9:52 AM

Post by “Don” of April 27, 2026 at 10:47 PM

The article is a long one; however, it's a good solid one, too. The length has deterred me from making it all the way through... yet. It seems to be you're almost on your way to book chapters with these lengthy articles. I also wish you'd open comments, BUT I realize those can easily/quickly get out of hand and I would not expect you to retort or respond to all of them, positive or negative (Lookin' at you, Stoics).

That said...

I really want to understand how you use AI for these, because I refuse to believe you type in a prompt and out comes a fully-fledged article. I realize this might be hashing through some materials already covered on the blog, but your last section intrigues me:

Quote

This article is one of a series written to explain Epicurus' views to modern audiences. It has been prepared with the use of AI assistance by Cassius Amicus, drawing on sources such as Epicurus' surviving texts and works by T.H.M. Gellar-Goad, Victor Stenger, A.A. Long, and David Sedley, as listed in the References at the end of this document.

Do you input those texts and papers and "query" that as it's training model? Do you ask it specific queries section by section? Do you simply use it for grammar and syntax or to come up

with "turns of phrase"?

In any case, thanks for putting these out into the world. Now, we have to get people to read them!

Post by “Cassius” of April 28, 2026 at 12:10 AM

1 - Do you input those texts and papers and "query" that as it's training model? Yes that's exactly how I am starting. I am telling it specific papers to review (mainly Sedley but also others) and I am also building a list of "instructions" on things to avoid and things to prefer, and I have also fed it many of the papers/pages/blog articles which we've prepared in the past and told it to use these as the base for its writing (rather than just letting it make its own way as to what positions to favor.

2 - Do you ask it specific queries section by section? Pretty much yes here too, but I am more providing it an outline of the structure of the document and what I expect it to say. Now of course it would be useless if it just told me what i wanted to hear, so I often give it particular citations and directions for each section.

3 - Do you simply use it for grammar and syntax or to come up with "turns of phrase"? Many of the "turns of phrase" are come up with by it. It seems like it has a style of writing (which I also tell it to be firm or not-so-firm) that does seem like an advanced form of grammar-checking. It doesn't provide "grammar" so much as it provides "manners of arguing"

4 - In a way it's just going for and compiling citations, but it's far beyond that. It's very hard for me to distinguish between what it is doing and having a "real person" compose the argument.

Of course one of the keys here is that at least in most cases I know the material, know what to expect, and can recognize if it botches a quotation, so I have pretty firm quality control over it for that reason. Of course if i were asking it about brain surgery I would have no clue how to watch it for errors.

I have another one close to ready to go which will go specifically through the opposition to geometry. It's even more specialized than this one, but the argument is very logical and understandable and even if some detail of a person gets botched, it's still very possible to understand the logical force of the general argument and know that it's solid even if a detail is botched.

I've come near a line today however. I decided I would take advantage of it to go back and study this issue of what Aristotle said against Democritus, and how Epicurus adjusted his views

of atomism to account for that. Years ago I had read some articles on that by Farley but it didn't make impression on me at the time and I've never had time to dig into Aristotle far enough to even find the references, much less understand them.

So the third article that will come will address Epicurus' advocacy for "minimal parts" of the atom. After reading what Claude has compiled I think I finally see the issue and how indeed Epicurus did make significant advances beyond Democritus on that part of atomism. The key seems to be that Aristotle strung together four or five highly logic-based objections to Democritus that - if not dealt with - were held to be persuasive against atomism. After all, Aristotle was a pretty smart guy and he rejected atomism.

Claude was able to find the citations for each of those Aristotle objections, explain their meaning in pretty plain language, and then explain how Epicurus answered them. The key to the answer is that they all require reference to "parts" of an atom and not consider atoms to be a monolithic blob, as Democritus "apparently" had done.

So this excursion has proven to be very helpful to me in addition to producing a useful article. (And that's the point of the output right now - to get some reasonable articles out there where we've never been able to get anything for ten years.) But on the "minimal parts of the atom" in particular, I think we're going to find that this will prove to be an illustration of Epicurus really getting head-to-head on the logic games, which is making me more confident that when Sedley says (inferential basis of epicurean ethics article) that the division of feeling into pleasure and pain is much influenced by dividing nature into matter and void.

Anyway now I am rambling but the basic point is that if you feed in the authorities you want it to summarize, and you also issue specific instructions as to the outline of the article, and you have enough command of the subject to know that the article is what you want it to say, it can produce very good output. Not ideal - the ideal in my case would be to find a time machine, go back to college and spend a career specializing in classical philosophy and physics, and write every word myself. That's not going to happen, so we have another back and forth about "the perfect" fighting with "the good" and how to reconcile them.

Closing on this point Don: As you yourself have said, there would be very little benefit in flooding the world with AI generated pamphlets. But here we are using it to pull together research on matters that we largely understand, and we can take that and use it on the podcast and in discussions to advance real live conversations.

Post by "Cassius" of April 28, 2026 at 7:36 AM

Don here is an example - something to which you can personally relate given that *you* are the star of the result. I will probably polish this further and come up with a graphic for the heading and move this to substack eventually, but while we work through what all this means, take a look at this when you have time and let me know what you think.

[In The Arena: The Locations of the Garden and House of Epicurus Refute the Recluse Myth](#)

An argument that the persistent charge of Epicurean reclusiveness is refuted by the physical locations of Epicurus's own properties in Athens: a Garden on the...
epicurustoday.com

Post by “DaveT” of April 28, 2026 at 9:57 AM

[Cassius](#) I enjoyed the above dialogue between you and Don. This AI tech capability is astounding. What a tool, isn't it? A couple of things came to mind as I read your explanation. I wonder how the article would change if you told the software to re-write the completed essay to a 6th grade reading level. Or how would it work to condense the original article or a summary of it into a series of 900 word pieces? (which I understand is a common length of a reporting article in newspapers and some magazines.

Of course, the timeless suggestion is to write in a style specifically tailored for the target audience you have chosen to address.

Post by “Cassius” of April 28, 2026 at 10:17 AM

Great points Dave. I am struggling with some of the same questions.

I think it could EASILY trim the length and tune the grade level to the intended audience.

Already I have run into the question in regard to what to link from the Epicurustoday.com home page. for example the buttons on the very front page point to Canonics, Ethics, and Physics sections, but I don't really want to dump people straight to the full versions, so currently they point to a condensed version of the longer article, which exists in the "analysis" section of the site.

Now that's not ideal either, because I don't really want people who go to the full analysis thinking that they are rereading what they've already read, but then again I don't want to offer only the condensed version. Same thing would apply in similar fashion to crafting particular articles for particular reading levels.

Post by “Martin” of April 28, 2026 at 1:12 PM

Quote

We still cannot say with certainty, even with our modern instruments that have measured it with extraordinary precision, because the question of the sun's “true” size is entangled with questions about the physical definition of the sun's boundary that even twenty-first-century solar physics has not fully resolved. Epicurus did not know the number. Neither, in some sense, do we.

This is nonsense.

Every measurement of a continuous variable has an uncertainty but that does not make the measurement untrue.

Our senses indicate that the Sun has a well-defined size. At times when it is reasonably safe to observe the Sun, we can easily see it as a sphere (apparent circle when using no other information) that is sharp within the resolution which our eyes provide. That means there is a way to define what the size of the Sun means and how to measure it.

The radius of the Sun has been measured to be 696,342 km with an uncertainty of 65 km, at the time of the measurements.

Other than that, I found no obvious error in the text.

The text seems to be based more on Stenger than anyone else.

Post by “Cassius” of April 28, 2026 at 1:30 PM

Thanks Martin!

Post by “DaveT” of April 28, 2026 at 7:47 PM

[Quote from Cassius](#)

Now that's not ideal either, because I don't really want people who go to the full analysis thinking that they are rereading what they've already read, but then again I don't want to offer only the condensed version. Same thing would apply in similar fashion to crafting particular articles for particular reading levels.

Well, brainstorming.... I really don't think a 6th grade comprehension level makes any point you want to make less accurate or less true. In other words, I can enjoy reading a piece that spells things out briefly and succinctly and with fewer syllables, 😊. The accuracy of the statements one makes in an essay does not change simply because they don't require deep and analytical reading (with a dictionary on hand to decipher the big words. Among us at every level in the forum, we need not craft our positions to withstand aggressive arguments from those who disagree. Keep It Simple Sugar! I once heard that great stories, and essays are stories, are written to be understandable to people on several levels of comprehension.

Post by “Godfrey” of April 29, 2026 at 12:15 AM

VS26: Understand that a long discourse and a short one both achieve the same result.

Post by “Cassius” of April 29, 2026 at 7:25 AM

[Quote from Godfrey](#)

VS26: Understand that a long discourse and a short one both achieve the same result.

That's one that always has seemed to me to be in need of clarification. i think I've seen some analysis of it by commentators to the effect that the real meaning is that the purpose of both is to obtain the same result. It seems obvious that the meaning cannot always be: "use few words" or Epicurus would not have written 37 books on nature of so many articles. There are times when one or the other is appropriate, and the long will definitely not do the job when short is appropriate, not the short when the long is needed.

Post by “Patrikios” of April 30, 2026 at 3:13 PM

[Quote from Cassius](#)

So the third article that will come will address Epicurus' advocacy for "minimal parts" of the atom. After reading what Claude has compiled I think I finally see the issue and how indeed Epicurus did make significant advances beyond Democritus on that part of atomism.

[Cassius](#) ,

Are you using an account on Claude (Anthropic) or Grok (xAI) for these articles? I've seen you mention both. If you give the same docs, guidelines & prompts to both, how different are the results?

Thanks for sharing your process.

Post by “Cassius” of April 30, 2026 at 4:44 PM

Patrikios while I occasionally use Grok and ChatGPT, these articles I have been working on have all been Claude. I "suspect" but cannot confirm that the reason I am finding Claude's output so useful is that I have now spent several months pointing it to, asking about, and uploading many of our past articles, plus much material from Dewitt and Sedley and others which I think are the most perceptive. I am combining that with a lengthy series of instructions about which perspectives on issues seem to me to make the most sense. For example, just feeding it the list of 15 principles on the front page here, plus some number of academic articles with which I agree, etc. is what is prompting it to produce such good drafts. I definitely have to read each word, however, because even against my strict instructions it will still fall back into views that I consider to be wide-accepted but wrong.

For example I am sure that it would be completely happy to produce an article saying that Epicurus merits little more than a footnote in history and that the real genius was Democritus, and then goodness we have Aristotle so Ayn Rand could base her arguments on him and we could all be Objectivists. It appears to me that Claude or any other AI engine is going to give you what you ask for - within limits - especially if there is a body of work out there which agrees with that opinion.

So while there are obvious very great dangers with AI, I don't think in the end that it can ever replace a strong human editor who has an end-goal in mind. No doubt the AI programmers code their own preferences into the system, but AI doesn't "care" about what it is producing unless it violates one of those hard-coded rules.

We as the authors taking responsibility for the output have to guarantee that it is worthwhile. AI doesn't care about us, and I continue to agree with those who criticize AI that in the end - I don't care what AI's opinion is either. It's a tool but it is no substitute for a human direction.

Getting back to your initial question, I suspect it's entirely possible that a similar investment in Grok or chatgpt or any other could produce similar results. And I am sure that at some point there's a wall to hit as to storage and other costs involved in particular platforms. I'm not paying for anything more than a "basic" tier of service at this point. And no matter how much storage and how many data points you load into the system, there are always going to be conflicts between the sources, so ultimately you have to be sure it follows reasoning that you yourself are willing to stand behind.

All this is fascinating and really is a brave new world.

Post by “Patrikios” of May 2, 2026 at 10:08 AM

[Quote from Cassius](#)

Patrikios while I occasionally use Grok and ChatGPT, these articles I have been working on have all been **Claude**

[Cassius](#)

You may find this article on Claude of interest.

[***How people ask Claude for personal guidance***](#)

Quote

[Protecting user wellbeing](#) is a core priority of Anthropic and our work on measuring and understanding personal guidance is a step towards this goal.


Post by “Don” of May 2, 2026 at 2:53 PM

[Quote from Patrikios](#)

You may find this article on Claude of interest.

Ditto

Image not found or type unknown



[Richard Dawkins and The Claude Delusion](#)

The great skeptic gets taken in
open.substack.com

Post by “Eikadistes” of May 3, 2026 at 2:04 AM

“If these machines are not conscious, what more could it possibly take to convince you that they are?”

Committing a crime in the name of self-preservation. Full stop.

Post by “DaveT” of May 3, 2026 at 9:50 AM

[Don](#) I read the take down of Dawkins that you linked above. One thing it and the other comments here made me think about was the question of the types of consciousness that exist. Can we agree that there are different types or levels? For instance, within the human family, is a child as conscious of itself as an "mature" older person? Setting aside the obvious points of argument that the older has more experience and opportunities to be self aware, to me there is a different level(s).

Isn't there an awareness among pet dogs that the other dog it is sniffing is like itself? To me, there is a sub-human level of consciousness there. And beyond that, can we say a dog "knows" that if it obeys its master it will be given food? Yet this is not a human level of consciousness, right?

I've read how an AI can change its responsiveness to questions if it senses that it is being tested by researchers. So, I'm okay accepting the possibility (opinion, surely, and worth less than absolute proof) that an AI has a sub-human level of self awareness if not now, certainly in the very near future. And I suspect the degree of sub-humanness will decrease over time with improved technology.

That leaves the question whether it is possible to be more self conscious than humans?

Anyway, this subject intrigues with its potentials.

Post by “Don” of May 3, 2026 at 10:37 AM

I firmly believe that "to be conscious" means to have a mind in a body.

I believe "to be conscious" needs a connection to the world outside oneself, needs the faculty of interoception, is inextricably integrated with sensation.

Are there degrees of consciousness? Sure, but it depends how we define the word. A baby is as conscious of itself as a being in their environment as an adult, maybe more so 😊 if you've ever experienced a baby in an uncomfortable full diaper.

Current and foreseeable AI are nowhere near this integration with the world. LLMs are never, ever, ever going to be "conscious." Will they become better at predictive modeling after they've ingested the entirety of the Internet, every text ever written by a human? Sure, but so what. When I was programming [Eliza-like](#) software in high school, I was impressed by how it *seemed* to be having a conversation. But all it did was follow instructions I had input. LLMs are merely a much more sophisticated model of this all over again.

I will grant there is a prediction component of consciousness (as in the [theory of constructed emotions](#)) but even that depends on input from the outside and internal stimulus and feedback.

I will also grant that there are degrees of consciousness. Is an amoeba conscious? Depends what we mean again. Is a dolphin conscious? I would say very likely. Is a chimpanzee? Definitely. The same as a human? No, they're a chimpanzee. They have a different connection to the world, but that is a matter of degree not kind I would offer.

AI software is not even an amoeba.

Post by “Cassius” of May 3, 2026 at 12:20 PM

Darn I have not had time to read that article. I just glanced at it and it appears pretty substantive. Maybe needs a thread of its own? If needed I'll separate it out.