I need to know! Discomfort in Machine Curiosity

Presented by Nadia M. Ady **Tuesday March 16th, 2021**

Click here to skip to the list of sources!

Transcript

This morning, my main character is the discomfort of curiosity.

Occasionally, I might allow the discomfort of curiosity to be upstaged by its host, curiosity, but make no mistake, in this tale, the **discomfort** of curiosity is the hero.

Now, it's a sort of funny thing. As a society, these days anyway, we love curiosity. You can find articles esteeming its value in business,[1] imploring concern for its disappearance in schoolchildren,[2,3] and even proclaiming it as the secret sauce to a healthy life.[4]

The idea that we might value something that's uncomfortable isn't too hard to see. Don't they always say that nothing that's good for you can be pleasant? But the hype around curiosity is relatively recent, and there were long periods of human history where curiosity was considered a vice, not a virtue.[5] And yet even over those periods, humans were choosing to expose themselves to situations that they knew perfectly well would make them curious. Today, we pick up crosswords, mystery novels, and even the latest bingeable series that we know will be permeated so thickly with cliffhangers that we are bound to lose sleep over it.

Yet, curiosity *is* uncomfortable.

When you first think of curiosity, you probably think of being curious about something you've never experienced. For example, a good storyteller can make us highly curious about what comes next. I'm certainly guilty of nights spent reading well into the wee hours because the idea of setting down my newest novel is so acutely painful that I just can't quite manage it. This is the discomfort of curiosity, where the idea of not finding out feels simply terrible.

---

[1] Gino, Francesca (2018, Sep). The business case for curiosity. *Harvard Business Review*, 48. https://hbr.org/2018/09/the-business-case-for-curiosity

[2] Berliner, Wendy (2020, Jan). 'Schools are killing curiosity': why we need to stop telling children to shut up and learn. *The Guardian.* https://www.theguardian.com/education/2020/jan/28/schools-killing-curiosity-learn

[3] Engel, Susan (2011, July) The Hungry Mind: The Origins of Curiosity. *WilliamsCollege YouTube Channel.* https://youtu.be/Wh4WAdw-oq8

[4] Heid, Markam (2020, Feb) Curiosity Is the Secret to a Happy Life. *Elemental*. https://elemental.medium.com/curiosity-is-the-secret-to-a-happy-life-3dc5d940d602

[5] Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, *116*(1), 75.

However, we can also be curious about things we think we know but can't quite bring to mind; you can surely think back to a time: Maybe you're watching something new on TV and spot someone who looks terribly familiar. What is the *name* of that actor? The way curiosity is typically understood in the literature, this tip-of-the-tongue phenomenon is usually included within curiosity.[6] I often find myself pausing my shows and using the internet to seek the answer, because continuing without knowing seems like it will make it impossible to focus on the show I want to enjoy.

In either case, the experience of curiosity, while temporary, can be achy or excruciating—uncomfortable. But I still watch television. I still pick up new books. People everywhere do, despite the clear risk of curiosity.

So, we voluntarily choose to put ourselves into situations that **do not feel good.** Why? Why would we do that? This is a question that really troubled early curiosity theorists. It seemed like a bit of a paradox.[7]

TIMESTAMP 2:21

In 1994, economist George Loewenstein provided an initial theory as to why we would voluntarily expose ourselves to the discomfort of curiosity.[8]

He suggested that because curiosity is uncomfortable, we appreciate the relief we feel when we satisfy curiosity. He pointed out that it's actually quite sensible to expose yourself to a curiosity-inducing situation if you expect the pleasure you'll experience satisfying your craving will compensate for the discomfort of curiosity itself.

Loewenstein is emphasizing a comparison between how we felt before and how we feel now. When we were curious, we felt discomfort; once we satisfied our curiosity, we don't feel that discomfort anymore: we feel better. But that comparison between how we felt before and how we felt after? It isn't the kind of comparison that animals really learn from. The prevalent theory is that animals *actually* learn from a comparison between how they expected to feel and how they actually felt.

To see how expectations affect animal learning, let's talk about dopamine! Dopamine is a brain chemical associated with pleasure, and with learning. Several of the parts of the brain responsible for dopamine are currently thought to guide learning by signalling the difference between how we expected to feel and how things really turned out, so we can update our mistaken expectations to more closely reflect our experiences!

---

[6] Metcalfe, J., Schwartz, B. L., & Bloom, P. A. (2017). The tip-of-the-tongue state and curiosity. *Cognitive Research: Principles and Implications*, *2*(1), 1-8.
[7] Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, *116*(1), 75.
[8] Ibid.

I'm going to try to make this more concrete by discussing some experiments that took place in the 90s, in the laboratory of the neuroscientist Wolfram Schultz.[9] In these experiments, the researchers recorded the activity of dopamine neurons in monkeys. This was a behavioural experiment: The researchers trained their monkeys to press a lever in response to seeing a light cue. If the monkey pressed the lever as they were supposed to, they earned a drop of apple juice. While the monkeys were first learning this task, if a monkey acted on the light cue, pressed the lever, and received their apple juice, then immediately after receiving their apple juice reward, the researchers found dopamine neurons would light up with lots of activity.

This figure shows the timeline of a trial of the experiment. Each of the little dots represents a single dopamine neuron firing. You can see that right after the monkey gets its apple juice, there's a short interval with heavier dotting than you see in the baseline activity during the rest of the trial. The same information is aggregated in the bar chart at the top.

Like I said, dopamine is associated with pleasure: monkeys like their apple juice. But this is only what happened in early trials. After some time spent training, the pattern of activity changed. I'm going to give you 20 seconds to predict for yourself what the dopamine neuron activity started to look like.

TIMESTAMP: 4:43

[Thinking Music for thinking time]

TIMESTAMP: 5:03

Okay, I don't want to let you hang too long. The answer is that the increased activity shifted. As the monkeys learned the task, the spike in activity shifted to occur just after the light cue. So, after training, what happens when the monkeys receive their reward? Nothing very exciting. So, after a spike in activity, dopamine neurons quickly return to a baseline level, and when the monkeys receive apple juice they had learned to expect, the activity stays at that baseline level.

Just as a side note: if you've heard of the classical conditioning experiments with Pavlov's dogs,[10] you're absolutely right to connect that to this experiment. Both Pavlov's experiments and Schultz's experiments were all about learning predictors. The monkey's response to the reward shifts to become a response to the light cue that predicts that reward as they learn.

Now, I'm going to give you a moment to make one more prediction. What do you think happened to the dopamine levels when a monkey who had learned this little task didn't get the apple juice they were expecting?

TIMESTAMP: 6:00

---

[9] Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593-1599.
[10] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press. 343-345

[Thinking Music for thinking time]

TIMESTAMP: 6:20

All right, here's the answer. That spike in activity we saw right after the light cue? It still happened. After all, when observing the light cue, the monkey had no idea things weren't going to go the same as in dozens of previous trials, so that same increase in activity makes sense. What's different is that right after the time they would normally have received their reward, there was actually a sudden **drop** in activity. You can see that there are fewer dots than the baseline level. The neuron activity reflected the difference from what they expected.

Just like those monkey's brains, our brains react when the world doesn't fit our expectations, and other behavioural evidence suggests that when the world doesn't fit their expectations is when animals really learn.

And curiosity seems to share some of its mechanisms with reward learning. In the last fifteen years, there have been several groups of neuroscientists that have induced and tested curiosity in human participants while looking at the participants' brain activity. And sure enough, dopamine neuron activity seems to be involved.

TIMESTAMP: 7:16

But curiosity is different from learning how you can obtain apple juice from the researchers that come and visit you every day. Once we have satisfied our curiosity, the satisfier isn't appealing anymore. If you were curious about what happened to the dopamine neuron activity when the monkey didn't receive its expected reward, you were probably glad that I gave you the answer a moment ago. Hearing me explain the answer was curiosity's satisfier. However, where the monkey would repeatedly enjoy apple juice, you're not going to be impressed with me if I repeat the answering chunk of my presentation over and over again.

If satisfying curiosity is simply rewarding, as George Loewenstein suggested, why don't we learn behaviour that leads us back to the same reward, as in, back to the situation that satisfied our curiosity? After all, that's what animals learn with food rewards.

It was this kind of conundrum that my supervisor and I were considering when we developed a simple model of curiosity that's really centred on the discomfort of curiosity and that I'd like to share with you today. Quick disclaimer: There's no evidence that our model describes the way things work in wet squishy animal brains. What we're most interested in is how it seems to reflect human curiosity better than models that use the simpler idea that whatever satisfies your curiosity is rewarding.

TIMESTAMP: 8:28

Our model of curiosity makes use of discomfort by considering what could happen if you expect the world to feel bad. We wanted to model discomfort like... the opposite of rose-coloured glasses. Where rose-coloured glasses make you see the world with optimism and as better than it might be, the discomfort of curiosity acts like prickle-glasses! Yes, prickles as in thorns. Your prickle glasses make any future you can imagine seem unpleasant and uncomfortable, *except* a situation that you think would satisfy your curiosity.

The natural thing to do, if the only bright point in your world is somewhere else within reach, is to take action to get there. Let me give you a concrete example. Let's say I'm on the Acorn TV streaming service, I've just started watching Agatha Christie's Partners in Crime,[11] and I've reached the end of the first episode. I don't want to give you any spoilers, but I'm sure you've seen shows of the kind, where, in the final minutes of the episode, the main characters, that the showmakers have worked hard to convince me to invest in over the first forty minutes, have gotten themselves into some trouble and I am dying to know what happens next. I'm sure you've heard others use the same turn of phrase, *dying* to know. All else equal, I am going to click play on the next episode, because I know that's where I'll find the knowledge I'm seeking, and in doing so, get myself out of this horribly uncomfortable state of not knowing. In this example, I am taking action to get myself to a state where I'm not experiencing the discomfort of curiosity.

TIMESTAMP: 9:46

Now, onto the slightly more complicated part. How is this model going to let us avoid a desire for one experience over and over again like a monkey wants to experience receiving apple juice over and over? Well, we're going to do this by making a quick switch. Instead of comparing realistic expectations of the world to what we experience, we're going to let the discomfort of curiosity infect our expectations, making our expectations much worse than it is reasonable to expect.

What will this switch cause us to learn? Earlier, I told you that animals update their expectations of how good situations are based on a comparison between what they expect to happen and what actually occurs. So what happens if we compare discomfort-tarnished expectations of the world with our neutral, and probably not-so-bad reality? The final difference is going to be quite positive. If you compare awful with neutral, things are better than you expect. So now, everything we experience while curious seems better than we expect, and we're going to let that unexpected goodness increase our preference for those experiences.

Let me just say that, with learning, we don't put all of our eggs in one basket so to speak and assume that a single occurrence of a reward is going to happen again in the same situation, rather, we build up confidence in a pattern gradually, over multiple occurrences. The second time the monkey sees the light cue, the monkey isn't going to particularly strongly associate it with the apple juice reward. It's going to take several more occurrences of the pattern. Similarly, if a pattern stops occurring, our expectation for it will fade away.

---

[11] Acorn TV. Agatha Christie's Partners in Crime. https://acorn.tv/partnersincrime/agathachristiespartnersincrime/

What does that mean for our model of curiosity? For things that don't happen multiple times while you're in a state of curiosity, including most of the things you experience and do to get to the situation that satisfies your curiosity, we are not going to see a preference build up—we're only seeing those things once! However, if there are experiences that repeatedly induce curiosity, this repetition can build up a glow of preference for those experiences, and to any repeated pathways to them.

With my TV show example, curiosity is going drive me to watch the beginning of episode 2, so I can see how Tommy and Tuppence, (the partners in crime mentioned in the title of the show) got themselves out of trouble. But watching it once is enough to relieve my curiosity, so there isn't enough repetition for me to develop a preference for watching the beginning of episode 2 anymore, even though that was what relieved my curiosity. But if I discover multiple curiosity-inducing shows, I'm going to develop a preference for any shared pathways I go through while curious, like flipping on the television and navigating to Acorn TV. And what's turning on Acorn TV likely to do? Put me into another curiosity-inducing situation!

Remember, this preference stems directly from the discomfort of curiosity. While curious, I'm going to give all my experiences a little boost in how much I'd like to experience them again. And that only happens because I expected things to feel worse than they were when I actually experienced them—discomfort!

I see this as a much more satisfactory view of voluntarily exposing ourselves to curiosity-inducing situations.

TIMESTAMP: 12:50

Time for a summary! In this model, the discomfort of curiosity gave us two things: first, it pushes us to take action to satisfy our curiosity. It pushes us to get out of our current situation and towards curiosity's satisfier. I am going to click play on the next episode. Second, it leads us to learn a preference for curiosity-inducing situations. The discomfort of curiosity gave us a preference for returning to those pathways that regularly result in curiosity being induced, as in, after experiencing several good curiosity-inducing shows on Acorn TV, my preference for turning on the television and switching to Acorn is going to have increased.

OK, fine, you know I like to watch TV, great. But is this a good thing?

I'm working in artificial intelligence. My colleagues studying biological intelligence in humans and other animals mostly get to study mechanisms of intelligence how they are. As in: What are the mechanisms and properties that exist in rats, cats, and psychology students today?

In artificial intelligence, on the other hand, we have the good fortune to be able to look at the work of those biological scientists and see not only the way those mechanisms are, but also see inspiration for the way mechanisms and properties *could be.* For me, this means, what *could*

curiosity be like? But rolled up in that is some responsibility to consider how those mechanisms *should be*. Since artificial intelligence researchers get to dream up something new, we have the ethical responsibility to consider not just the *is* but the *ought to be*.

That said, I want to consider whether we want those properties that the discomfort of curiosity gives us. Should I be trying to implement machine learning systems with those properties? So here's one of the exciting things about human curiosity that makes me say yes.

There's been some recent work by Shirlene Wade and Celeste Kidd that suggests curiosity might actually point us towards fruitful learning experiences.[12] Our curiosity might act as a signal to indicate knowledge that we're ready to learn.

I think this is really exciting, the idea that curiosity might help us scaffold our own learning. However, it leaves us with an important open problem: what kind of mechanisms and algorithms will best allow us to estimate and recognize when we are ready to learn something, so we can make these mechanisms part of our machine curiosity methods?

If we can answer that question, then the two properties provided by the discomfort of curiosity are properties that we want. By pushing us to satisfy our curiosity and by giving us a learned preference for curiosity-inducing situations, the discomfort of curiosity makes us expand our knowledge effectively. And no matter how uncomfortable it is to be dying to know at the time, I'm grateful for these mechanisms that help me learn.

[12] Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic bulletin & review*, *26*(4), 1377-1387.

List of Sources

9:04 Acorn TV. Agatha Christie's Partners in Crime. https://acorn.tv/partnersincrime/

0:18 Berliner, Wendy (2020, Jan). 'Schools are killing curiosity': why we need to stop telling children to shut up and learn. *The Guardian.* https://www.theguardian.com/education/2020/jan/28/schools-killing-curiosity-learn

0:18 Engel, Susan (2011, July) The Hungry Mind: The Origins of Curiosity. *WilliamsCollege YouTube Channel.* https://youtu.be/Wh4WAdw-oq8

0:18 Gino, Francesca (2018, Sep). The business case for curiosity. *Harvard Business Review*, 48. https://hbr.org/2018/09/the-business-case-for-curiosity

0:18 Heid, Markam (2020) Curiosity Is the Secret to a Happy Life. *Elemental*. https://elemental.medium.com/curiosity-is-the-secret-to-a-happy-life-3dc5d940d602

0:33 Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, *116*(1), 75.

1:38 Metcalfe, J., Schwartz, B. L., & Bloom, P. A. (2017). The tip-of-the-tongue state and curiosity. *Cognitive Research: Principles and Implications*, *2*(1), 1-8.

3:30 Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593-1599.

5:33 Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press. 343-345

14:15 Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic bulletin & review*, *26*(4), 1377-1387.