Is Human Behavior Just Elaborate Running and Tumbling?

Nathaniel Travis
contact@nathanieltravis.com

January 17, 2022

Abstract

Here we present the *reward-taxis* model of human nature. It's both a formal model and an intuitive explanation of many aspects of behavior. The model is based on recent work by Karin and Alon demonstrating the equivalence of the dopamine circuit in animals and the chemotaxis circuit in bacteria. The dopamine circuit performs what they call reward-taxis navigation. Here we extend that work, arguing that humans use reward-taxis in a fundamental way and that reward-taxis explains many psychological phenomena.

Keywords: *reward-taxis; human algorithm; growth imperative; irrational behaviors.*

1 Introduction

You know that scene in the movie, *A Beautiful Mind*, where John Nash is studying pigeons “hoping to extract an algorithm to define their movement”? Well, it appears we’ve found that algorithm, and it defines the movement of humans too. Furthermore, it provides us with a new lens for understanding psychological phenomena.

In getting to this algorithm, one key insight is to realize that the dopamine circuit—which simultaneously controls movement, motivation, and encoding reward prediction errors in humans and other vertebrates—serves the same role as the chemotaxis circuit in bacteria, and is algorithmically equivalent as well.

This insight came to me merely insofar as it led me to google “dopamine chemotaxis.” Upon googling, though, I found that Karin and Alon had taken the insight much further to building a formal model and demonstrating its experimental accuracy (Karin & Alon, 2021). They began with the following puzzle, which we can pose again here: why is it that the dopamine circuit simultaneously controls movement, motivation, and reward error encoding? In other words, what do all of those things have in common?

The answer is that the dopamine circuit is playing a *navigational* role akin to chemotaxis (*chemo* meaning *chemical* + *taxis* meaning *movement*). Just as bacteria use a scale-invariant run-and-tumble algorithm to navigate towards chemical attractants (rewards), humans use a similar scale-invariant run-and-tumble algorithm to navigate towards more general rewards. The difference is primarily in how we determine our expected rewards. While bacteria determine their expected rewards based on a passive sensing of the environment, humans determine our expected rewards based on a complex, predictive model of the world. And, of course, we have different definitions of what a reward is.

Perhaps it's not obvious that such a realization will lead to psychological insights, but as we'll see, it does in fact provide the foundation for a rigorous and unifying understanding of many psychological phenomena. We'll need to establish a few things before we get there, though, so first, let's take a look at bacteria.

2 Bacterial Chemotaxis

Bacteria use a process known as chemotaxis to move towards chemical attractants, such as food, and away from chemical repellents, such as toxins. It's an incredibly well-studied process. As you can imagine, back in the 1600s when cells were first observed under a microscope (by Leeuwenhoek and others), one of the first questions they probably asked was "how are these cells moving?" The modern biochemical study of chemotaxis was pioneered in the 1960s (e.g. J. Adler, 1966). The modern mathematical and algorithmic study of chemotaxis was pioneered in the 1990s (e.g. Barkai & Leibler, 1997; Alon, Surette, Barkai, & Leibler, 1999). And today it's understood down to the minute biochemical and algorithmic details.

The process (or algorithm) consists of a biased random walk of "runs" and "tumbles." During a run, the bacterium coordinates its flagella to swim in a straight line. During a tumble, one flagellum kicks out, causing the bacteria to spin and start swimming in a different direction. The process is random (a random walk), but it's nevertheless effective in finding attractants due to being biased towards attractants and away from repellents.

The way the biasing works is the following: the bacterium keeps track of the change in the chemical gradient as it runs. If it is climbing up the gradient, it will tend to keep running in that direction and not tumble since it "smells" that it's getting closer to the attractant. The tendency of running vs tumbling is controlled by a single parameter, the tumbling frequency ϕ (tumbles/sec). If the bacterium is climbing up the chemical gradient, tumbling frequency goes down. If it is going down the chemical gradient, tumbling frequency goes up. It's quite simple. But despite its simplicity, this run-and-tumble algorithm guarantees that, on average, the bacterium will head in the right direction and obtain its rewards.

Figure 1: Runs and tumbles

There is one other crucial aspect of this algorithm, which is how tumbling frequency changes as a function of the change in the chemical attractant level. If the attractant level goes up, tumbling frequency goes down; however, if the attractant level ceases to *continue* going up, tumbling frequency will return to its *steady state* value. In other words, it's adaptive.

Specifically, it is adaptive in a way that is known as *scale invariance*, which means, for example, that the attractant level going from 1 to 2 causes the same effect as going from 2 to 4 (or 1000 to 2000). Another way of describing scale invariance is that changes are always relative: the absolute values don't matter; only the relative difference compared to the baseline matters. Furthermore, this scale invariance has the property of *exact adaptation*, which means that tumbling frequency levels will return to their steady state when there ceases to be any change in the chemical gradient. This exact adaptation is evolutionarily crucial: if it were not exactly adaptive, it would not be robust to changes in the environment and likely would not be able to survive and reproduce long-term (we'll see shortly how animal reward-taxis has this same property of exact adaptation).

The scale invariance is implemented in bacteria with a simple feedback loop: the protein that increases the rate of tumbling also decreases its own rate of production, so it's a self-regulating system that is designed to always return to its steady state.

Figure 2: Exact adaptation

The formal chemotaxis model, as well as pseudo-code, can be found in the appendix. But for now, our intuitive understanding should suffice to begin to understand the “animal algorithm,” which we’ll see is strikingly similar.

3 Animal Reward-taxis

Animal movement has a superficial similarity with bacterial movements in the sense that, well, animals also run and change directions. However, Karin and Alon (K&A) demonstrated that this relationship is not merely superficial; animals are using the exact same algorithm.

K&A called this process in animals reward-taxis (*reward + taxis* meaning *movement*), since animals have a broader definition of reward than simply chemical attractants/repellents (e.g. food, play, shelter, mates, etc.). If animals are in an open environment—such as mice running in an open field or zebrafish swimming in open water—their movement looks like a series of runs-and-tumbles up the *expected reward* gradient, and it also exhibits the property of scale invariance.

Specifically, the K&A model translates the bacterial chemotaxis model to the animal reward-taxis model via the following changes:

1. The 2D (or 3D) map of chemical levels c is translated to a 2D (or 3D) map of expected

rewards R

2. Tumbling frequency ϕ is translated to average run duration ϕ (seconds)
3. Change in the chemical level Δc is translated to change in expected reward ΔR
4. When a “tumble” occurs, rather than sampling a new direction from a uniformly random distribution, we sample according to the distribution of expected rewards

In the same way that chemotaxis is scale invariant with respect to chemical attractant levels, reward-taxis is scale invariant with respect to expected reward.

What exactly are “expected rewards”? Well, as we mentioned, all animals have a notion of rewards. And these rewards correspond to specific, numeric values, which can be measured via dopamine responses in the brain. These numeric values, however, are always *relative* (hence the scale invariance property), and in fact, dopamine response roughly corresponds to the temporal derivative of expected reward (recently, this was shown by precisely manipulating mice’s reward expectations via virtual reality (Kim et al., 2020)). The “expected” part of “expected rewards” means that animals are always making a prediction about future rewards.

Actually, chemotaxis itself is a special case of reward-taxis. In bacteria’s case, expectations are based simply on sensing the environment and keeping track of the change in chemical levels (and when they tumble, they lose all sense of expectation, hence why they have uniformly random tumbles).

Figure 3: Animal runs and tumbles (adapted from Karin & Alon, 2021)

The experimental evidence for the K&A model is that it explains the *generalized matching law of operant behavior* (Baum, 1974; Herrnstein, 1961). Furthermore, it explains many physiological studies on the relative nature of dopamine, and provides an elegant, mechanistic explanation of dopamine's threefold functionality (movement, motivation, and reward error encoding).

Thus, to summarize, K&A demonstrated that animals, such as mice and zebrafish, navigate up expected reward landscapes using a reward-taxis algorithm akin to bacteria's chemotaxis. Which leads us, naturally, to the question of humans.

4 Human Reward-taxis

Humans are animals, and humans do reward-taxis just like animals. However, the runs-and-tumbles of people are not as visibly obvious as the runs-and-tumbles of animals in an open environment.

4.1 Hierarchical Runs-and-Tumbles

In the case of human reward-taxis, it's possible to think of a run as *striving towards a goal* rather than physically running in a straight line. This is a stretch beyond what K&A showed, but we can make the equivalence between movement and striving towards a goal more clear by considering an example: pursuing a career in X. Such a goal is likely a big, long-term goal; however, progress towards that goal actually consists of a hierarchy of progressively smaller runs-and-tumbles:

1. In order to pursue your career in X, you need to get a job doing X. You may "run" towards a particular job A, only to be rejected and forced to "run" to a different job B.
2. Furthermore, in order to get job A or job B, you need to go to the interview (let's say you have to drive there). Hopefully, you'd be able to follow one driving route to get there, but in the off-chance that a road is blocked or you hit heavy traffic, you may need to take a different route.
3. Furthermore, to get the car to drive to your interview site, you may try to use the car you own. Likely it will work, but in the off-chance that your car won't start, you may need to call a taxi. Etcetera.

Figure 4: Hierarchical runs-and-tumbles

At the base level of these hierarchies, striving does always involve some spatial movement, even if it is merely internal. Thus, we can begin to see how—besides a different perception of expected reward—human’s runs-and-tumbles are analogous to the runs-and-tumbles of animals in an open environment.

4.2 Parallel Runs and Partial Tumbles

Another peculiarity of human reward-taxis (likely including other animals), is that we perform multiple runs-and-tumbles in parallel and, to some extent, simultaneously. For example, we can walk to the grocery store to pick up food for dinner (pursuing the reward of food—and possibly other rewards if the dinner includes other people), while simultaneously imagining solutions to a scientific problem (pursuing career or intellectual rewards). Furthermore, we can do *partial tumbles*. For example, we can switch jobs while maintaining the same living situation, or switch living situations while maintaining the same job. We can also do more drastic tumbles, changing our direction in several dimensions at once.

4.3 The Feeling of Exact Adaptation

Furthermore, If we consider the *feeling* of reward expectations to correspond to *desire*, then the idea of exact adaptation (i.e. scale invariance) of expected rewards should sound familiar: just when we feel like we’ve gotten what we want, we want more. An extreme example of human’s exact adaptation was given by Dan Gilbert: his team tracked lottery winners as well as paraplegics. Naturally, the lottery winners were ecstatic at first, while the paraplegics were distraught. But only a year later, their subjective well-being levels were equal (Gilbert, 2009). The psychology literature is littered with other examples of this phenomenon (e.g. Furnham & Argyle, 1998; Easterlin, 1968). In other words, we—like all animals—have a *growth imperative*, a need to continually climb the expected reward gradient.

4.4 The Feeling of Run-and-Tumble

To give more intuition for why we are in reality using a run-and-tumble algorithm, consider that in life when we’re “running” in a consistent direction—for example, in a job, in a relationship, or in solving a problem—so long as we’re seeing small improvements, we’ll tend to continue in the same direction, even if that direction is not optimal. It’s only when our growth slows or halts, that we “tumble” and change directions. In other words, our run duration ϕ does appear to be a function of how our reward expectations change over time ΔR . And this is not merely intuition; research backs up this notion. For example, the number one reason people switch jobs is because they are concerned about the lack of opportunities for career advancement (LinkedIn, 2015). It’s all about growth.

4.5 What Space Are We Navigating?

So far, we’ve been arguing that we humans do a form of reward-taxis navigation. However, it’s clear that we are not simply navigating the 3D spatial world, moving up and down, left and right, like bacteria or mice in an open field. The question is, then: what space are we navigating?

The answer, to put it simply, is that we’re navigating the space of all our potential actions. These actions include spatial movement. They also include things such as imagination or logical deduction (mental processes), as well as language communication (which involves some amount of physical movement and mental processing).

There is a hierarchical and chain-reaction nature to actions. At the base level, every action is a combination of physical movement and mental processing; however, these base actions can be organized into a complex set that becomes, for all practical purposes, a single action (for example, driving a car or calling a friend on the phone). To complicate the idea of what constitutes an action further, consider that navigating our expected reward landscape of all possible actions can lead us to new possible actions, and many rewards

are, in fact, a means to unlocking more possible actions. In other words, the space we're navigating is dynamic. Suffice it to say, though, that the space our reward-taxi is navigating is the *action space*.

4.6 The Human Algorithm

Putting all this information together (and taking the complex notion of expected reward as a given for now), we can present a simplified, discrete-time view of the "human algorithm":

Inputs: initial objective environment E , position within the environment x , set of N possible actions $A = \{a_0, a_1, \dots, a_N\}$, subjective N -dimensional expected reward map M , scalar expected reward level R , change in expected reward $\Delta R = 0$, run duration set to its steady state value $\phi = \phi_{st}$, run direction d

Reward-taxi:

while *alive* **do**

$new\_R \leftarrow$ maybe obtain reward and update expected reward

$new\_ \Delta R \leftarrow new\_R - R$

$\Delta R \leftarrow (1-w)*new\_ \Delta R + w*\Delta R$ (ΔR is a running mean of change in expected rewards)

$R \leftarrow new\_R$

$\phi \leftarrow f(\Delta R, \phi, \phi_{st})$ (update run duration 'proportionally' to ΔR)

$heads \leftarrow$ flip coin with probability $p(\phi)$ of being heads (p decreasing with ϕ)

if *heads* **then**

$d \leftarrow$ sample new direction according to M ('tumble')

$actions \leftarrow$ take a step in direction d (i.e. do the actions corresponding to d)

$E \leftarrow$ update environment (based on *actions*)

$x \leftarrow$ update position (based on *actions*)

$A \leftarrow$ update possible actions (based on our new env, pos, and learnings)

$M \leftarrow$ update expected reward map (based on new env, pos, learnings, and actions)

done

There are, of course, a lot of "implementation details" here, including the fact that this algorithm is running in a highly distributed, decentralized fashion (humans are not simple, merely using a top-level algorithm that is, at its core, simple).

4.7 Quick Clarifications

Before moving on, let's clarify a couple of important questions:

1. K&A proved their results for the midbrain dopamine circuit. What about the forebrain?

It is a leap to take results shown for the midbrain dopamine circuit in zebrafish and mice and extend it to humans with a distinctively large forebrain. However, given our understanding of the function of the forebrain—for example, its role in prediction and reasoning (Bubic, Von Cramon, & Schubotz, 2010)—it is reasonable to assume that its role is not to change the “outer loop” of the reward-taxis algorithm but merely to modulate the expected reward term inside. At the very least, if we take this assumption as an axiom for now, we'll see that it will lead us to a useful framework.

2. What are the “nice” properties of reward-taxis that make it so universal? Are there any alternative algorithms we use or could use?

We did mention a “nice” property of scale invariance, namely, exact adaptation. However, we have not discussed the nice properties of run-and-tumble. One such nice property is that it is incredibly flexible. By changing the dynamic expected reward term, any complex behavior can be elicited. One other nice property is that it works well as a *distributed, decentralized algorithm*. This property can be observed in bacterial chemotaxis: each of a bacterium's flagella are operated by independent motors, which are connected to independent sensings of the environment. If one single flagellum “decides” to tumble, the whole bacterium will tumble. Yet when running, they are all able to work in concert.

This distributed, decentralized aspect could be beneficial for humans as well. Consider that we often detect something urgent from one of our senses—for example, feeling a burst of pain when touching a hot surface, or seeing an object flying at us. Having a distributed response system means that that one sense can directly boost dopamine levels, causing us to move and escape the danger. If our biological system was not distributed and each sense or body part had to directly communicate with each other, our response times would likely be much slower.

That being said, we could imagine alternative algorithms, for example, instead of “randomly” running-and-tumbling, what if we did a kind of breadth-first search, systematically exploring the environment? One possible reason we might not do this is the fact that exploration is a learning process: we often don't have a good enough mental map of the environment to be able to systematically explore. However, it's important to also note that run-and-tumble is not actually random in the case of reward-taxis: directions leading to a higher expected reward are more likely to be chosen. Furthermore, run-and-tumble is flexible enough to implement a systematic exploration (just imagine the rewards dynamically changing to create a pseudo-breadth-first search). Thus, there may not exist alternative

algorithms that provide additional benefits.

4.8 Expanding the Expected Reward Term

Now that we've established the outer loop of the reward-taxi algorithm, the question becomes, what is going on with the expected reward term? This map is where all our human "magic" is happening, where the value of our consciousness and predictive model of the world must lie. Reward-taxi itself doesn't explain the details of the expected reward term; however, we can make some base assumptions that will be helpful in our discussion of psychological phenomena.

Assuming for simplicity that our actions are deterministic, in other words, if you take action a in state s_t , then the next state $s_{t+1}(a, s_t)$ is deterministic, then our numerical value of our expected rewards adheres to a Bellman equation like this (Bellman, 1954):

$$R(s_t) = r(s_t) + \gamma R(s_{t+1}) \quad (1)$$

In each state we get an immediate reward r (though r can be 0, i.e. no reward), followed by all our future rewards (which are discounted by a constant factor γ since immediate rewards are more useful than future rewards).

We could break this equation down by viewing expected reward as a sum over all possible reward outcomes. However, a more useful way to look at it in order to gain psychological insights is to view it as a sum over different *reward types*. For example, money, accomplishment, love, food, friends, influence, sex, power, etc.:

$$R(s_t) = R_{mon}(s_t) + R_{acc}(s_t) + R_{love}(s_t) + R_{food}(s_t) + \dots \quad (2)$$

Furthermore, if we focus on "big" rewards (i.e. rewards with a large magnitude that we don't get often—for example, working hard in order to get a large year-end bonus k time steps in the future), we can rewrite the equation in terms of paths to these big rewards:

$$R(s_t) = \gamma^k r_{bonus}(s_{t+k}) + \gamma^{k+1} R(s_{t+k+1}) + (r(s_t) + \dots + r(s_{t+k-1})) \quad (3)$$

Intuitively, we can think of γ^k as a distance-term (distance to the reward you're seeking) and $r(s_t) + \dots + r(s_{t+k-1})$ as a path-term (these notions of distances and paths will help us to explain some psychological phenomena):

$$R(s_t) = dist(r_{bonus}, s_t)(r_{bonus}(s_{t+k}) + \gamma R(s_{t+k+1})) + path(r_{bonus}, s_t) \quad (4)$$

5 A New View of Psychological Phenomena

5.1 New Tools

Now that we've established an understanding of reward-taxis, let's see what new tools it gives us for an analysis of human nature and psychology:

1. A visual and geometrical view of ourselves as an agent navigating an expected reward landscape
2. A graphical and mathematical understanding of the scale invariant nature of desire and our "growth imperative"
3. An algorithmic view which lends itself to algorithmic analysis (e.g. what are the extreme cases for the algorithm? What do those extreme cases correspond to in real life?)
4. The idea that all action is optimizing for expected reward and everything else is essentially a "modifier" of expected reward

The rest of this paper will briefly touch on several aspects of psychology, conjecturing on how our new tools and perspective can provide new understandings of psychological phenomena.

5.2 Quirks and "Irrational" Behaviors

In the reward-taxis view of human nature, all actions can be seen as rational given an expected reward landscape (though, this expected reward landscape differs from person to person). This can help us to understand the rationality behind behaviors that appear "irrational" economically or evolutionarily. For example, consider the following:

1. Why do some "gym rats" build their muscles well past the point of being a way to attract potential mates?
2. Why are some people willing to free solo climb icy mountains despite the danger?
3. Why are some people willing to have fewer children in order to pursue career goals?
4. Why are most billionaire tech CEOs interested in space travel and extraterrestrials?

There are no doubt potential evolutionary explanations for each of these phenomena. But, from a purely evolutionary standpoint, these phenomena are at least (a) not obviously explained or (b) seemingly irrational.

In the reward-taxis framework, however, each of these can be explained simply as a byproduct of our growth imperative, our need to continually climb the expected reward gradient. In reward-taxis, desire for mates and reproduction can be seen as one of many

reward types. Certainly, evolution has made it a large, important term in the reward equation (and certainly, reward-taxi itself is an evolved mechanism), but the desire for different reward types is subjective and subject to change. What is common to all, though, is our desire to keep growing in some subjective direction.

5.3 Autonomy

Next, let's discuss "autonomy," an important concept in many psychological theories. In reward-taxi, there are many ways to understand autonomy. Reduced autonomy can be understood as an increasing of the distance to many desired reward types (equation 4). Viewing ourselves as navigating agents, we can imagine restricted autonomy as an artificial warping of our reward landscapes that restricts our movement.

This also gives us a new view of the concept of intrinsic vs extrinsic motivation and "motivational crowding out" (Ryan & Deci, 2000). In the reward-taxi model, what is important is not a distinction between intrinsic vs extrinsic but (a) whether we're able to continually sense that our expected rewards are growing, (b) whether we "tumble" and change directions, and (c) whether distances to our rewards are too high.

Thus, if we look at some classic studies of motivational crowding out—for example, taking kids that are intrinsically motivated to draw, then giving them extrinsic rewards for drawing, and subsequently taking those rewards away to find that the kids no longer appear to be intrinsically motivated to draw (Lepper, Greene, & Nisbett, 1973)—we can interpret these studies differently than they've been historically interpreted. For example, when the additional extrinsic reward is taken away, the kids will no longer feel that they're climbing the reward gradient and thus are more likely to tumble to a different activity. It has nothing to do with the fact that the reward was extrinsic, merely that the kid's relative growth was stifled.

5.4 Depression

Reward-taxi does not immediately tell us what the source of depression is, but it does lend itself to conjecture. For example, there are at least two intuitive ways depression can be understood in reward-taxi. The first way is our exactly-adaptive desire system being "broken" in the sense that it is not exactly adapting as it normally does. Specifically, being stuck in a lower-than-usual steady state. The second way is being in a *local optimum* of the expected reward landscape, a state where whatever action you take, you see a decrease in expected reward. We could call this second way our reward expectations being "broken." Both ways can lead to a state of constant tumbling (or slower movement if we consider K&A's alternative model where dopamine regulates running speed).

How, though, could our desire system possibly get stuck in a lower-than-usual steady state? Isn't it designed to be perfectly robust and adaptive?

Well, it's important to note that, in reality, biological systems that exhibit exact adaptation are not exactly adaptive over an infinite range of inputs. If input levels are sufficiently high or sufficiently low, the systems may break (for example, in the case of chemotaxis, it's adaptive over five orders of magnitude of inputs). Thus, dopamine levels becoming sufficiently high—for example, via chronic stress—could be a cause of problems.

There is a more realistic hypothesis, though. If we consider the fact that dopamine is a relative measure of expected rewards, we could hypothesize that in depression we're stuck constantly comparing our current expectations to some alternative state, for example, constantly ruminating over a past state, or constantly dreaming of an alternative state that were restricted from.

5.5 Personality

Interestingly, even bacteria have different personalities. And even genetically identical bacteria have different personalities. Specifically, some bacteria tumble much more often than others (and some run much longer on average than others). This is due to the fact that although tumbling frequency always returns to its steady state, different bacteria have different steady state values (Spudich & Koshland, 1976). This can be seen as analogous to introversion vs extraversion in humans, which is driven by none other than dopamine (our equivalent to the tumbling frequency protein in bacteria) (DeYoung, 2013).

Beyond introversion and extraversion, we can see how other, more complex personalities could emerge from reward-taxis dynamics. Specifically, looking at the distance term in expected reward from equation (4), $dist(r, s_t)$, we could say that the distance to a particular reward always depends on some "skill." Also, using a particular skill to seek a reward will develop the skill further, reducing the distance to other rewards that rely on that skill. Thus, it creates a feedback loop of relying on and developing a particular skill more and more. A simple case of this dynamic can be seen as the cause of right-handedness vs left-handedness, and in fact, it's also the mechanism Jung proposed as the driving force behind personality differences (Jung, 1921/2016). In that case, the "skills" correspond to discrete, specialized sets of modules in the brain, and preference for certain skills over others manifests as personality differences.

5.6 Motivation and Needs

Many psychological theories—for example, those of McClelland, Deci & Ryan, and Maslow, among others (McClelland, 1987; Ryan & Deci, 2000; Maslow, 1943)—describe human motivation via a set of categories describing our needs, drives, and/or motivators. These categories are insightful and often heavily experimentally validated, however, they do not necessarily represent core causal factors of motivation (and experiments do not typically test whether they are core causal factors).

Thus, reward-taxis, being a mechanistic explanation of our underlying motivation system can be a helpful aid to these existing psychological theories of needs, drives, and/or motivators. Furthermore, reward-taxis can help us understand how and why different categories of needs are different.

For example, let's consider Maslow's dichotomy of growth needs vs deficiency needs. Many studies have questioned whether this dichotomy is a true dichotomy, or whether there is a hierarchy of needs as Maslow claimed (Wahba & Bridwell, 1976). Certainly some needs are absolutely required for survival, and others are important for good health. But there is another side to Maslow's distinction, which we can understand in reward-taxis as the *shape of reward functions*.

Some reward types are quickly satiated at a particular threshold, and beyond that threshold, an additional amount of that reward is not even perceived as rewarding. These can be thought of as the deficiency needs (or deficiency rewards). Other reward types are unbounded, never fully satisfied, or reach diminishing returns more slowly. These are the growth needs (or growth rewards). There is a spectrum of such reward functions, which can be incorporated in our above equations by making rewards a function of what how much of that reward we already have (which can be captured in the state s_t).

5.7 Repressed Needs/Desires

We can also use reward-taxis to re-interpret Freud's idea of repressed needs/desires and sublimation. Freud claimed that desires that are frowned upon by society, such as certain sexual desires, could be "sublimated" to acts of higher social value, such as scientific, artist, or career pursuits (Freud, 1930/2015). In the reward-taxis model, this is due to the fact that although some reward types (e.g. sex) may have large magnitude, they may also have high distance (e.g. due to social constraints), thus causing other reward types with lower magnitude but shorter distance to correspond to a higher overall expectation (and thus causing our adaptive desire system to compel us towards these other reward types).

5.8 Consciousness

Lastly, let's briefly discuss consciousness. Reward-taxis does not explain consciousness, but it does support the idea that the *purpose* of consciousness is to help us in building a predictive model of the world (in order to make more accurate predictions about future rewards).

6 Conclusion, Limitations, and Future Directions

Here we presented the *reward-taxis* model of human nature. We claim that reward-taxis does not merely describe certain narrow aspects of human behavior but that (1) all action is based on expected reward and that (2) the algorithm used to act based on expected reward is the reward-taxis algorithm, which appears to be shared, in some form, by all organisms from bacteria to humans.

That said, there are many limitations to this model. For example, understanding how we determine our maps of expected reward is likely more important than understanding the reward-taxis algorithm that utilizes expected reward (the details of these “maps” and how they’re learned over time is essentially what separates bacteria from humans). Although we were able to make some high level assumptions about reward expectation, reward-taxis itself does not explain it.

Another major limitation is that this is a single-agent model. Any interaction between people would have to be captured within the expected reward, which gives us a limited view, especially given that we’re such a social animal. In other words, it’s primarily viewing humans from a third person perspective and, to some extent, a first person perspective, but it’s missing the “group perspective.”

Besides those limitations, there are many areas that are not necessarily limitations but which we were not able to get into, such as the following (which could be good future directions):

- Incorporating learning, uncertainty/risk, and energy expenditure into the model
- Going deeper into particular areas of psychology
- Validating reward-taxis in humans by seeing if it predicts behavior (in some well-defined scenarios)
- Considering applications to AI
- Looking deeper at the “nice properties” of the reward-taxis algorithm (why does it appear to be shared by all life forms?)
- Considering how to “hack” the algorithm (i.e. live better lives)

References

- Adler, J. (1966). Chemotaxis in bacteria: Motile escherichia coli migrate in bands that are influenced by oxygen and organic nutrients. *Science*, 153(3737), 708–716.
- Adler, M., & Alon, U. (2018). Fold-change detection in biological systems. *Current Opinion in Systems Biology*, 8, 81–89.
- Alon, U., Surette, M. G., Barkai, N., & Leibler, S. (1999). Robustness in bacterial chemotaxis. *Nature*, 397(6715), 168–171.

- Barkai, N., & Leibler, S. (1997). Robustness in simple biochemical networks. *Nature*, 387(6636), 913–917.
- Baum, W. M. (1974). On two types of deviation from the matching law: bias and under-matching 1. *Journal of the experimental analysis of behavior*, 22(1), 231–242.
- Bellman, R. (1954). The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60(6), 503–515.
- Bubic, A., Von Cramon, D. Y., & Schubotz, R. I. (2010). Prediction, cognition and the brain. *Frontiers in human neuroscience*, 4, 25.
- DeYoung, C. G. (2013). The neuromodulator of exploration: A unifying theory of the role of dopamine in personality. *Frontiers in human neuroscience*, 762.
- Easterlin, R. A. (1968). The american baby boom in historical perspective. In *Population, labor force, and long swings in economic growth: the american experience* (pp. 77–110). NBER.
- Freud, S. (1930/2015). *Civilization and its discontents*. Broadview Press.
- Furnham, A., & Argyle, M. (1998). *The psychology of money*. Psychology Press.
- Gilbert, D. (2009). *Stumbling on happiness*. Vintage Canada.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the experimental analysis of behavior*, 4(3), 267.
- Jung, C. (1921/2016). *Psychological types*. Routledge.
- Karin, O., & Alon, U. (2021). The dopamine circuit as a reward-taxis navigation system. *bioRxiv*.
- Kim, H. R., Malik, A. N., Mikhael, J. G., Bech, P., Tsutsui-Kimura, I., Sun, F., ... others (2020). A unified framework for dopamine signals across timescales. *Cell*, 183(6), 1600–1616.
- Lepper, M. R., Greene, D., & Nisbett, R. E. (1973). Undermining children's intrinsic interest with extrinsic reward: A test of the "overjustification" hypothesis. *Journal of Personality and social Psychology*, 28(1), 129.
- LinkedIn. (2015). *Job switchers survey* (Tech. Rep.). Retrieved from <https://business.linkedin.com/content/dam/business/talent-solutions/global/en-us/job-switchers/PDF/job-switchers-global-report-english.pdf>
- Maslow, A. H. (1943). A theory of human motivation. *Psychological review*, 50(4), 370.
- McClelland, D. C. (1987). *Human motivation*. CUP Archive.
- Ryan, R. M., & Deci, E. L. (2000). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology*, 25(1), 54–67.
- Spudich, J. L., & Koshland, D. E. (1976). Non-genetic individuality: chance in the single cell. *Nature*, 262(5568), 467–471.
- Wahba, M. A., & Bridwell, L. G. (1976). Maslow reconsidered: A review of research on the need hierarchy theory. *Organizational behavior and human performance*, 15(2), 212–240.

Appendix

A1. Chemotaxis Model

In the language of systems biology, the chemotaxis circuit implements a *non-linear integral feedback loop*, which has the following model (M. Adler & Alon, 2018):

u is input (chemical from outside the bacterial cell), x is an internal node (a protein), and y is the output (a protein). ρ is a constant (the ratio of removal rates of x and y). The dot notation means the change/derivative (e.g. dx/dt).

Here is pseudo-code for the chemotaxis algorithm. For a simplified view, we imagine that time is discrete:

Inputs: initial environment E , position within the environment x , attractant level C , change in attractant level $\Delta C = 0$, tumble frequency set to its steady state value $\phi = \phi_{st}$, run direction d

Chemotaxis:

while *alive* **do**

$new\_C \leftarrow$ measure attractant level

$new\_ΔC \leftarrow new\_C - C$

$ΔC \leftarrow (1-w)*new\_ΔC + w*ΔC$ ($ΔC$ is a running mean of change in attractant level)

$C \leftarrow new\_C$

$\phi \leftarrow f(-ΔC, \phi, \phi_{st})$ (update tumble frequency 'proportionally' to $-ΔC$)

$heads \leftarrow$ flip coin with probability ϕ of being heads

```
if heads then  
     $d \leftarrow$  sample new direction uniformly (tumble)  
  
     $x \leftarrow$  move one step in direction  $d$   
done
```