The Will Of The Free

After all the vessels of philosophy have been coaxed into the harbor, after all the boats carrying all the theories of mind & their mechanisms have been tied snugly to the piers of science, one vessel seems always to remain stubbornly at sea—and painted on its bow are two words: *Free Will*.

Do humans possess free will? There are a plethora of ways to approach this question, but most can be divided into two categories: philosophical & scientific. We might think of the philosophical approach as attempting to answer the question: *what truly is free will?* In contrast to that broader query, the scientific approach attempts to discern the specific relationship between the neural activity that determines our actions and the neural activity that generates our conscious contemplation of those actions. In other words, science attempts to answer the chicken or egg question: *which comes first, thinking about doing it or deciding to do it?*

In this last essay, we'll use the currents from both of these approaches to finally bring our wayward vessel, *Free Will*, into harbor alongside the rest of Narrative Complexity's multi-faceted fleet. But we'll begin by exploring that more concrete question about the neural relationship between thinking & deciding; this is the question that's most directly addressed by the elements & mechanisms of our theory here. And since free will is essentially about whether we are consciously capable of "choosing" our actions, at the center of this question is the relationship between our *conscious mind* & our *unconscious (or subconscious) mind*.

One Brain, Two Minds

How, then, do we define these two entities of mind? In the view of Narrative Complexity, our *conscious mind* is essentially that *experience* of hearing the internal dialogue generated by our loop of consciousness. The locus of this part of our loop (our *Dynamic Core* of conscious awareness) appears to be located within humans' highly-advanced prefrontal cortex.

Thus, our actual consciously-experienced mind is merely a *portion* of that loop—the narrative parcel that finally arrives on our conscious highway of internal dialogue & becomes one of those candidates for a short- or long-term memory. This experience includes our "awareness" of the current

moment (or of those surrounding elements of the moment that our conscious focus has prioritized & is actively attending to).

Some of the sensory, internal & emotional data from that conscious "awareness" can be attached to those narrative parcels as they are seared in our storage, creating a more full memory of that experience—a memory that is a combination of internal dialogue narrative parcels & associated spike experiential data that occurred simultaneously. Thus—when we recall those right-brain-based narrative memories —what we are recalling was originally (in its *very first* iteration) founded upon a low-fidelity, attention-defined recording of our *conscious mind* in those moments.

When we compare this view of the conscious mind with the many mechanisms in our system that appear to occur in the loop *well before* that conscious dialogue actually arrives in our heads, we can get a good idea of just *how much* of our consciousness-generating mechanisms are actually occurring within the realm that we would define as the *unconscious mind*. In fact, you might even say that Narrative Complexity is as much a theory of *unconsciousness* as it is a theory of consciousness.

In truth, thus far we have left the majority of the unconscious mind out of the discussion, and have mostly focused on those elements of mind that contribute directly to our conscious experience—like emotions & internal dialogue. Luckily for us, there's no need for our theory to present its own detailed explanation of the mechanisms of the unconscious mind, because in 2008, Yale University's John A. Bargh & Ezequiel Morsella published a new & compelling view of the unconscious mind[1] whose perspective aligns perfectly with our theory's approach to the matter. (And Morsella's own insightful *Passive Frame Theory*[2] presents a useful framework for understanding why the *conscious* arena has evolved specifically to help direct the kinds of voluntary, adaptive skeletal-muscular actions that we'll be discussing.)

The basic thrust of Bargh & Morsella's argument is that the currently dominant view of the unconscious—essentially equating it with the "subliminal"—provides far too narrow a perspective on its genuinely robust systems. In many ways, most current discussions of our subconscious treat it as an undercurrent of thought that subtly & unknowingly influences our much more dominant conscious faculties. Bargh & Morsella argue that in reality, our unconscious mind is the primary (& ancient) machinery of the brain, and our conscious systems are actually a final layer of experience set atop that machinery. This is a perspective that Narrative Complexity wholeheartedly shares. (It's also shared by Gerald Edelman, the author of the *Dynamic Core Hypothesis*[3] that our theory relies heavily on.)

This Old Hominin

Here's one way to think of it: roll back the evolutionary clock to the time before humans developed self-aware consciousness —back to those brains that had not yet developed the capacity for internal dialogue. Since we know that even the most advanced apes do not appear to possess the full capacities of human internal dialogue, we might assume that those virgin branches of our hominin ancestry likely arrived on the evolutionary scene without those full capacities having yet been developed. (An assumption we explored back in essay #1 while discussing our *Terrence Deacon-supported* [4] theory of language evolution.)

How might we think of these ancient hominin, post-ape minds? In essence, these brains would have primarily consisted of that integrated, multi-sensory, Dynamic-Core-based arena of experience (*aka*, the consciousness viewfinder), plus all of the mechanisms that we now consider to be the *unconscious mind.* At this point in mammalian brain development, creatures were already cross-coordinating & analyzing highly-sophisticated sensory input, using the cerebral cortex to store that incoming data in a large array of complex, modular & associative hippocampus-defined memories, and using that cortex to apply inborn & learned rules to emergent current & pinged-via-association stored data in order to help construct unique & dynamic behavioral responses—a cognitive loop that also produced neurotransmitter-fueled proto-emotions and was managed by organs like the thalamus & basal ganglia, *and* that used the amygdala to aid in emotional production/regulation & for ancillary storage of intense, primal pain- & fear-based memories. (This loop is depicted by our *Quick Sketch of Pre-Language Mammalian "Cognition"* in the Appendices.)

These are, indeed, the primary mechanisms lurking beneath & helping to sustain *modern* human consciousness. This old original hominin system is merely lacking the ability to "talk to itself" about what it's doing. Basically, the most elemental aspect of mammalian consciousness—the Dynamic-Core-based arena of experience—is a part of *all* vertebrate brains (going back to those *lampreys*). But until mammals, this conscious arena is only used to direct entirely reflexive & pre-programmed (aka, robotic) motor responses.

With the emergence of modularized experiential data patterns in mammalian brains, that conscious arena began to help direct dynamic & creative *learned* motor responses, allowing mammals to diversify their behavior based on their environment & experiences. And this kind of modular-but-non-linguistic mammalian cognition is exactly what advanced apes & early hominins employed. Humans have merely (& spectacularly) added a language-based inner voice that allows us to internally

contemplate & manipulate that experiential data & those motor responses—aka, *self-awareness*.

I believe, in fact, that all of our mammalian (& probably all of our earlier vertebrate) ancestors did (& do) experience a type of non-language-based consciousness that is in many, many ways very similar to our own experience of consciousness. Yet, we are unable to truly fathom what the difference is between these experiences, because we simply cannot place our minds in a world in which there is *not* a word for *everything* (you can't even contemplate the *notion of it* without using words to contemplate it).

Ultimately, the complexity with which we are able to contemplate that huge amount of nuance, detail & interactivity in the world around us via our internal dialogue's system of language represents a *massive leap* in the nature of the conscious experience that is produced. And this perspective actually presents us with a pair of seemingly-contradictory truths. On the one hand, our language-based internal dialogue clearly sets the nature of our consciousness apart from the experience & capabilities of all other earthly animals. On the other hand, the essential mechanisms that define consciousness exist in a rather fluid continuum across vertebrate (& advanced *cephalopod*) species—and most animals do, indeed, experience a kind of consciousness that has much more *in common* with our own than it has differences.

How To Behave Like A Human

What does this really mean in terms of experience & behavior? It means if you kick a dog, the dog will feel (*experience*) pain, and sadness or anger, and it will know & remember that you (& the kick) are the source of its pain, sadness or anger, and it will both express this feeling in its behavior toward you and attempt to take action to protect itself from further pain—pain that its brain is cognitively *predicting* might happen based on both this & previously-stored experiences. That's some pretty heavy-duty lifting on the part of this little canine's primitive, limited consciousness.

And yet, the dog cannot *think* to itself or try to specifically *communicate* to you: "That makes me sad & angry because that both physically hurts & because I thought you liked me, and when you kick me it makes me feel like you don't like me, and I like you & I really want you to like me. And let me tell you why I like you..." The dog simply *behaves* in all of these ways (& experiences these feelings) based on those dynamic, modular cognitive systems that are making use of all those sensory input tools, stored data, motor scripts, switchboards, and behaviorally-governing, neurotransmitter-fueled emotional mechanics. (Of course, being humans, upon observing this dog behavior we are prone to naturally *interpret* it as the dog expressing that complex dialogue.) Those aforementioned mechanisms are all part of that essentially non-*thinking*

system of pre-human consciousness that does everything but talk to itself.

This kind of brain & the resultant behavior are much like the "zombie" human that we contemplated in our first essay—a less creatively- & dynamically-efficient version of a human that does not possess a capacity for internal dialogue. Imagining that zombie helped us to demonstrate why it was evolutionarily advantageous for hominins to eventually morph this capacity for modular cognition into an ability to dynamically "tag" those modular elements with words & transform our behavioral rule system into a syntactic rule system—language—that could make use of those words to help generate much more creative & robust predictions & solutions.

When we look at the human mind this way —as a machine that is capable at its evolutionary core of operating in an almost fully-responsive capacity without employing internal dialogue—it is clear that the *vast majority* of our behaviors & actions can be controlled entirely via our *unconscious* mind. For example, think about reading an article and occasionally reaching for & drinking from a coffee mug while reading. In cases like this, your consciousness is almost fully-engaged in your primary task: the article's narrative. But when you glance at the mug or experience an internal pang of "coffee-desire" your attention might be diverted enough to momentarily think "coffee" between the other sentences in your head.

This data input (the sight or the pang) triggers both the thought "coffee" & related motor scripts that enact your physical coffee-drinking routine. This motor script is partly enacted because there is no conflicting unconscious inhibitory script also being triggered—like pain in your free hand that negates the act of reaching and might then force you to pause reading & consciously free your other hand from the magazine (or iPad) to reach for the mug.

If the coffee-drinking motor script is uninhibited & immediately enacted, your consciousness almost-simultaneously returns to the article. You're not thinking to yourself "grab the mug, bring it to my lips, sip"—at this point in your life, drinking coffee in this circumstance is a deeply-rote script and can be enacted without that kind of focused conscious attention. That's because, according to our theory, once a deeply-rote, ongoing motor script like this is triggered, it can be maintained through a spatially- & physically-informed *unconscious* "action loop" that is routed to motor areas *without* engaging conscious cognitive processes. (This "action-maintenance loop" performs the same function in pre-language mammals, who can also use their more primitive cognitive processes to trigger dynamic motor responses that don't require ongoing cognitive direction, but do require ongoing physical/spacial maintenance.)

Thus, as you drink the coffee—although your own unconscious act of drinking is immediately within your realm of "awareness," allowing you to peripherally perceive these coffee-consuming actions as you read—your own act of drinking may or may not become a *true* part of that conscious & remembered experience. This is why, immediately after finishing the coffee (which might've come as a complete surprise) if asked how many times you actually drank from the mug or at which points in the article you drank, you might have no idea—even though you were vaguely aware of every act of drinking that took place while you read.

And when we look closely at our lives & the actions that compose our days, it appears that most of it is actually a result of our unconscious mind humming away like it has in mammals for eons: responding to incoming data with a barrage of subconscious competing/cooperating motor scripts that have no need for internal dialogue in order to function and maintain sophisticated behaviors & actions. Thus, to actually behave like a human (as opposed to any other animal) is not really to *do* any of the things that we do—it's to do them, and *while* doing them, think something like: *"Man, that was stupid."*

Motor Task Chunking

Apologies for the *academic-paper-ish-ness* of our section heading here, but the poet in me (which is, oddly, how I began my brain expedition) could not resist the strange, thick music of the words: *Motor Task Chunking*. And the words do, indeed, describe exactly what we're going to discuss: the chunking of the consciously-attended-to "grab the mug, bring it to my lips, sip" into the almost-entirely-unconscious *coffee*.

The best place to begin exploring *Motor Task Chunking* (okay, I'll stop) is way back in those pre-language mammals (like dogs & monkeys) who use modularly-constructed proto-narratives in the process of recording high-priority experiences & building dynamic cognitive responses. In those early minds, in order for the modular components of this cognitive process to trigger actual actions (which is the whole point) those components would have to link to or trigger specific & appropriate motor scripts. (The thalamus' & basal ganglia's switchboards aid in transmitting & processing those often-competing cognitive & purely-reflexive motor scripts en route to the motor cortexes that coordinate their execution—something we'll discuss in more detail later.)

In humans, this cognitive process is overtaken (& powerfully, exponentially enhanced) by complex language. Nonetheless, the same fundamental relationship exists (as it *must*) between those narrative components & motor scripts. Thus, in humans, the words of our internal dialogue can (and often *do*) lead directly to actions. (Although—to give an awfully revealing *sneak preview* of our

ultimate verdict on free will—the motor instructions that result from those words are likely triggered by the just-generated dialogue in the micro-moment *before* the dialogue's appearance within our Dynamic Core allows those words to be *heard* by us.)

In this system, the more complex & elaborate the motor script you can tie to a single word or thought, the more efficiently your consciousness can "off-load" the handling of full, multi-step motor sequences to those motor systems (sequences that our consciousness doesn't *really* need to be involved with). When we first learn a complex sequence, those full, multi-step motor scripts simply don't exist yet. Thus, we need to cognitively break that sequence into the smaller components for which motor scripts already exist.[5]

This means we have to actually *think* the words that trigger those smaller, already-learned components: *grab the mug*, *bring it to my lips*, *sip*. This is obviously a fairly inefficient way to drink something—and a pretty criminal misuse of those magnificently creative & robust systems of human consciousness. Our brain is *very* interested in getting a full, multi-step motor sequence in place for this mundane (but still very necessary) action so it can trigger it via that *blink-of-an-eye* "coffee"—and thus, keep its valuable conscious energy focused on more rewarding matters (like the *in*valuable insights coming forth from Sarah & Vinnie on the *97.3 FM Morning Show* that you listen to as you drink that coffee).

When we're *hoping* to trigger one of those fluid, multi-step motor scripts (usually one learned through intense practice) but *instead* end up *thinking* ourselves back into those less-fluid un-chunked component scripts, it has a specific & feared name: *choking*. This is why it's so bad when a Major League second baseman starts to freak out and *think* about the actual physical act of throwing the ball to first base. Suddenly that fluid, multi-step motor script is being interfered with by those cognitively-expressed smaller, component (& less fluid) motor scripts that are now actively *conflicting* with the more fluid one.

In an act as precise as throwing a baseball at the velocity that a Major Leaguer does, any tiny hitch or *oh-so-momentary* conflict in what muscle is doing what (because you suddenly can't stop *thinking* about the various aspects of throwing that baseball) can have professionally-disastrous results. This is also exactly why the *anxiety* produced by *worrying* about whether or not you can make the throw leads to the same problem that likely seeded the doubt (*all too often: one bad throw leads to another*).

Back in essay #2, we explained how those pain-based (or *survival*) emotions are designed to make our actions & thinking less fluid in favor of a more "hyper-aware" state. (Essentially, a decrease in focused, fluid attention is the cost of an overall lower-risk & diffusely-attentive state.) When we're nervous or anxiety-ridden, our brains are predicting a bad result from the

upcoming action or event, and thus, want us to *slow down & think about this! And don't just stare at the problem—look around & see if you can come up with something better!* These are not the ideal neural conditions for throwing out the runner at first.

And this kind of problem is mostly a case of our *conscious mind* sticking its nose in where it doesn't belong—and where it *claimed* to supposedly have no interest, which is why we built the multi-step motor script in the first place. But our consciousness is like the ultimate *helicopter parent*—and as soon as it suspects that you're about to make a *valuable mistake* (the anxiety tattled on you), it has a tendency to step in and try to assert its ever-guiding influence over the whole matter. Although this can occasionally lead to some embarrassing scenes in front of your friends (or 40,000 baseball fans)—more often than not, when it really counts, the *hesitation* being counseled by our consciousness is *exactly* what the situation calls for.

When this whole process is working efficiently, however, it can allow someone like an experienced pitcher to calmly survey the batter, devoting all of his conscious mental capacities to the many nuances of this momentary conflict between them. Because he has learned & developed the highly-fluid & elaborate motor scripts necessary to enact widely-varied versions of throwing the ball, he can trigger a series of complicated actions via the tiniest part of a cognitive thought (& some closely-related, suddenly-applicable complex script might even step into the process almost *unconsciously*—like reflexively zipping the ball to first base when he ever-so-slightly detects the runner leaning a little too heavily toward second).

With his cognitive processes freely devoted to his engagement with the batter, the pitcher can bring all of that additional data to bear on the execution of the complicated, unconscious & fluid motor scripts. Instead of worrying about the actual act of pitching, he's calculating the nuances of the entire pitcher-batter conflict. It's the difference between a Little League pitcher & an experienced hurler. Without all of that practice in turning smaller component motor scripts into multi-step scripts—the complicated act of throwing different pitches to different batters in different game situations requires so much conscious cognitive work that there's simply *no room at the neural inn* for calculating anything like nuance.

So, besides how the words *sound*, that's what's so great (& occasionally vexing) about *Motor Task Chunking*.

The Diffuse Box Of Consciousness

In essence, this continually-dynamic conscious experience layered atop a deep, robust, ever-percolating well of unconscious activity is what Narrative

Complexity refers to as the *Diffuse Box of Consciousness* (a concept we've touched on briefly in the previous essays). I use the term "diffuse box" because although the experience of human consciousness can be specifically defined as one part of our internal dialogue loop, the content of that dialogue is continually & powerfully shaped by all of the data that is also being processed & managed within those same loops, but does not necessarily ever emerge in our internal dialogue.

Keep in mind (pun *sort-of-intended*) that the final singular narrative parcel we actually *hear* is the result of associating, comparing & culling *myriad* data patterns that have been pinged by the just-consumed linguistic, environmental & physical input. This means that even though you may not have *consciously noted* the painting of the sailboat at the dinner party (i.e., you were only peripherally processing that visual data while talking to the attention-consuming attractive attendee) you still *unintentionally* end up steering the conversation in a slightly different direction because you once almost *drowned* in a sailboat accident (due to how that peripheral-but-powerful data input impacts your unconscious data & dialogue *culling process*).

This might seem like a pretty sneaky & possibly harmful way to shape behavior, but it makes sense that our brains would seek to *slightly weight* our cognitive processes in one direction or another based on peripheral *currently-low-priority-but-historically-significant* data within our immediate environment. To make matters even more fluid, consider: even if the *first round* of processing didn't allow the data to escape the unconscious, the subsequently-pinged data processed within our loop can easily gain enough prestige or attention to allow that original unconscious thought's or action's *next generation* to emerge in soon-to-be internal dialogue. (i.e., Suddenly interrupting your conversation with that attractive attendee with an apparent non-sequitur like *"That sailboat painting is kind of freaking me out."*) Thus, consciousness is best described as something that both has distinct borders & a highly-fluid, interwoven-with-the-unconscious nature: a *diffuse box of consciousness*.

In a strange way, living within that box and inhabiting most of its *diffuseness* is the old hominin within us, dumb but not *dumb*—silent, but still capable of learning & enacting nearly *anything* that you teach or show or even ever-so-momentarily *request* of them. And floating atop the box—concretely & securely, but *barely-just-above* the diffuseness—is us, language-based human consciousness: a transformation of the most prestigious emergent diffuseness into a fluid, flexible, highly-integrated thread of thought & experience, a dynamically & complexly organized narrative equation born to be consumed again by the unconscious & help seed all the thoughts to

come. We are the effervescent result of all the diffuseness percolating beneath us.

Imagine the loop of consciousness perpetually cycling, the majority of its data only either sparking other unconscious data via association or triggering primarily unconscious motor scripts. Based on the focus of our current actions, emotions, behavior, environment & attention, a tiny sliver of the highest-priority data running that loop ends up taking the exclusive & primarily single-narrative-only route into the arena of our conscious awareness, emerging in the form of *self-heard* internal dialogue—a thread of language-based thought woven into a complex Dynamic Core of experience. That route also includes a round-trip back to the unconscious processing in which it originated, where it rejoins the great unwashed masses of recycled thoughts & incoming data—some of which thanklessly results in those myriad unconscious actions, associations, emotions & behaviors.

From this perspective, we might think of the evolution toward modern human consciousness as a neural infrastructure initiative that slowly built both a new kind of high-tech engine to go around the same old track, and a new *extension* of the track that only accommodates this high-tech engine—a track that takes the engine to & from an exquisitely-evolved & high-tech consciousness-inducing destination within our prefrontal cortex.

Of course, the question we're really seeking to answer here: did that new engine bring with it *true* free will?

The Sliding Scale of Story, The Executive & The Virtuoso Switcher

Before we delve into the soul-defining answers to that question, let's explore a closely-related narrative mechanic (nostalgically taking one last brief *detour* in an essay-to-essay journey that has already taken generous liberties in indulging detours). This deeply-applied narrative mechanic might actually *only* be considered via detour—because its powerful influence on decision-making is hidden away in the very nature by which we tend to construct & consider the narratives that surround the decisions produced by our loop of consciousness.

This mechanic relates to how we instinctively (usually due to previous emotionally-biased—& thus, *recurring*—cognitive patterns or "habits of thought") frame the *scale* of these decision-generating narratives using our more advanced human predictive skills. We'll cut this detour right to the chase by providing a common example: the decision to photograph one's child as the child walks toward the school entrance on their first day of kindergarten. For our purposes, let's imagine this scenario confronting a naive first-timer dropping off their only child—thus, we will assume that this parent has no previously-

stored (personal or observed) experiences upon which to base a more reflexive, habitual response to the moment. First, consider how an earlier version of the mammalian brain might approach this decision (forgetting for a moment that dogs & such don't even know *how* to take photos). That kind of dialogue-less system would be likely to frame this choice within a more narrow narrative scale that is ultimately unconcerned with the moments that extend beyond the very near future.

But humans—mostly due to those language-based, boundary-pushing predictive & future-imagining capabilities—are able to project from the current first-day-of-school moment to a future moment in which they will want to remember & re-experience (& *not lose forever*) the *emotion* of this moment. According to our theory, this *self-projection* into the future is ultimately another one of those predictive cognitive tasks, and thus relies on those linguistically-based tools of internal dialogue. Because our first-time parent has no previous similar (thus likely to be pinged) experiences, they will need to consciously consider that they might want a photo later in order to arrive at the choice to take the photo now. Since it is not a rote & automatically triggered motor script under these conditions, the act *cannot* be triggered without generating the desire & possibility through internal dialogue contemplation. This act is a unique, new choice in an unfamiliar environment & circumstance, and thus is likely to rely heavily on the creative-choice-generating capacities that our internal dialogue mechanisms specialize in. Without contemplating & articulating (internally or verbally) the future desire for the photo, there is no reason in the current moment to actually take the photo—since (if we imagine our scenario in that *real-rolls-of-film* pre-Facebook era) it provides no in-the-moment pleasure nor satisfies any other in-the-moment need. This is why almost no non-human brain would even *consider* taking the photo—it can't *imagine* any reason to.

Within most non-human brains (except probably mammals like dolphins, elephants & advanced apes) that *in-mind* future version of the animal—and most of the imaginable possible moments that might occur between now & then—don't *really* exist to them or are even *available* to their minds for genuine creative & manipulative contemplation. Those other mammalian brains *do* use prediction patterns to help calculate & select most-beneficial responses based on considering predicted near-future results. Nonetheless—even if the predicted result requires multiple steps to achieve—this is really an extrapolation of symbolic & associative causal logic, not a true creatively-conjured & *mentally-malleable* internal self-depiction of the future.

By using this kind of sliding scale of story, our brain essentially views every decision as a "hinge" along a narrative continuum—preceding the hinge is the narrative history leading to this story moment, and following the hinge is the predicted future of this narrative. In making a decision, the brain seeks to inhibit or activate this hinge in a way that is most likely to result in a desired future narrative (or to avoid an undesired one). Thus, how we internally conjure, perceive & scale both past *and* future elements of the current narrative impact how the brain chooses to flip the decision-hinge.

This makes the scaling of story an *essential part* of most decisions that we consciously contemplate via internal dialogue. For example, if a cigarette smoker views the act of smoking within the smaller-scale narrative of the stress of the previous/current moments and the pleasure they currently desire/will soon receive, the smoker will very likely choose to light the cigarette. But if the smoker in that moment expands—via internal dialogue—the scale of that narrative to include the likely future prediction of themselves dying young from lung cancer, they might be more hesitant to light the cigarette in that instance (or be quick to add a new twist to their predicted future narrative, like imagining *quitting next week* in order to inhibit the inhibition to smoke produced by the thought of dying young). Of course, in the latter case, even a smoker whose brain doesn't bother to add a new future-twist might very well light the cigarette without hesitation, which is part of the problem with smoking & its sinful kin: they produce some *very* powerful desires for their temporary pleasures. Which conveniently brings us to our next stop along this detour: the battle between urges & narratively-reasoned desires.

Back in essay #2 we noted that those urges (mostly based on pure physical desires or threat responses—aka, our human versions of those ancient "proto-emotions") tend, at their very highest levels, to outweigh most opposing narratively-generated desires in our decision-making process. Whenever we are starving/parched, deep in the throes of lust, completely exhausted, totally repulsed by disgustingness, or in the grip of fight/flight (*or seriously craving a smoke...or really, really angry*) the desire to satisfy that powerful urge will often influence our action-choice more strongly than any narratively-generated (and/or belief-based) impulse to act otherwise.

The ultimate results of these kinds of consciously-contemplated decision-making battles appear to be primarily determined by our “executive control” area, the *dorsolateral prefrontal cortex*[6]. Other areas might be involved in helping the DLPFC to mediate this process, but it’s generally viewed as the decision-making hub. (*And these battles can also be impacted by that endorphin-based willpower discussed in essay #2*). Although an urge (or a *craving)* can originate via unconscious sources, once we

become consciously aware of the urge, it can be contemplated via internal dialogue —allowing our choice to *inhibit or encourage* the urge to be handled by that executive control area (which our narrative building/analyzing mechanisms feed data into). And in the view of Narrative Complexity, our language-based parcels of internal dialogue (built by our left hemisphere narrative-building mechanisms) must be "processed" in two key ways *before* the DLPFC is able to use those narrative parcels when deciding to inhibit or encourage an urge or action.

One: the dialogue needs to generate—likely via categorical, associative (& habit-driven) methods—a specifically-selected (potential) action trigger (or a set of *cooperative triggers*) that might aid in satisfying our urge or narratively-based desire. Thus, this associatively-based *potential-action selection process* can use an (anger-induced) thought like "I'm gonna kill him!" to help select a goal-specific (& *emotion-satisfying*) potential-action trigger among multiple options—like reaching out to strangle the target or seeking out the nearest heavy object.

The anger (& the supporting threat-response proto-emotion *fight/flight*) that *induces* & helps to define the *intent* of this potential-action-triggering thought is *initially* produced by the previous (& likely observational) thought—something like, "He's trying to hurt my child!" (*therefore*, "I'm gonna kill him!"). And if we really want to achieve our *malevolent intent*, our chosen action trigger had better be capable of cuing a well-practiced, elaborately-chunked, complex & highly-fluid motor script.

In addition to spurring this *potential-action selection process*, every just-built parcel of dialogue must also be analyzed according to all of those myriad *emotional & proto-emotional equations* (which judge our 13 primary emotional pairs & various physical urges). As explained in Essay #2, these "equations" make a wide range of judgements about factors like value gain/loss, prediction success/failure & belief compliance/violation (all of which are involved when analyzing a narrative parcel like "I'm gonna kill him!"). Consequently, this process of emotional & urge analysis involves a wide range of brain areas, such as (but not limited to) the orbitofrontal cortex, the anterior cingulate cortex, the amygdala, and the insula—all of which could process this narrative/emotional data in parallel (although it wouldn't be surprising to find interaction between *some* of these areas during processing). And this emotional analysis occurs at the same time as that *potential-action selection*.

Therefore, we hypothesize that the DLPFC receives our language-based narrative parcel (with its newly-attached & just-selected potential-action trigger or triggers) in conjunction with this wide variety of inhibiting & encouraging emotional judgments produced by that wide array of emotion-analyzing brain areas. The executively-controlling DLPFC is then able to compare & calculate separate-but-relevant emotional judgements—such as

predicted value gain/loss (excitement/fear), prediction success/failure (confidence/anxiety) and belief compliance/violation (satisfaction/guilt)—allowing all of the varied emotional factors to contribute to our decision to encourage or inhibit any narrative parcel's potential-action trigger(s).

Our model of emotion's role in decision-making is similar to the model that Oxford neuroscientist Edmund Rolls proposes in his 2014 book, *Emotions and Decision-Making Explained* [7]. One key difference, however, is in the placement of action selection processes. Although Rolls' model places action selection *after* emotional analysis, in our view, all of the necessary data required to select an appropriate action trigger (i.e., intent, action goals, action targets & action obstacles) is present in the syntactic & vocabulary content. Thus, because this is the same data that's used for emotional analysis—and because emotional judgements do not seem necessary in *selecting* a potential action—it seems more likely & efficient for action selection & emotional analysis to occur simultaneously.

In addition, from our perspective, if the emotional judgements that ultimately determine action activation/inhibition also contribute to the action-selection process, then our brain would already have all the data necessary to decide activation/inhibition at the time of choosing the action. This means there wouldn't technically be any need to ever emotionally *inhibit* any potential action choice, because we *already know* what we're emotionally willing or not willing to do when we choose our action. In such a scenario, why would our brains ever choose an action that it *already knows* it isn't emotionally willing to enact? This would also seem to make the DLPFC less of an executive and more of an *executive assistant*.

In our model, the same raw data—syntactic & vocabulary content—is processed by our emotional analysis & action selection mechanisms simultaneously, and the results of both are reunited in the DLPFC to determine whether the action is activated or inhibited. Basically, if the emotional judgements that support or encourage any urge- or narratively-based possible-action(s) neurally *outweigh* competing emotional judgements that discourage or inhibit the potential-action(s), then the selected action trigger(s) are "chosen" to be enacted (or vice-versa, preventing or inhibiting the selected action). Ultimately, from our theory's perspective, this decision-making process is not usually a traditionally-conjured "this-action-or-that-action" choice. Rather, it is a choice between enacting or inhibiting a singularly-directed action (or set of actions) that satisfy a specific urge or narrative desire.

In other words—returning to our homicidal example—when someone else's anger-inducing act sparks the impulsive thought "I'm gonna kill him," we aren't choosing

between strangling him or, say, merely spitting on him. The language-based narrative intent spurs the selection of an appropriate urge/narratively-satisfying action trigger, e.g., reaching out to strangle him. The DLPFC's job is to then use the proper emotional judgements to choose to enact or inhibit that singularly-directed action (or set of actions). Thus, we are choosing to strangle him or not to strangle him in that moment, not choosing between strangling him or taking some other specific-but-non-lethal action. Of course, because our consciousness is a looping *speed demon*, we can inhibit our desire to strangle & follow it up with an uninhibited desire to *expectorate* so quickly that it basically & perceivably *feels* like we're choosing between the two (unpleasant) options in the same moment.

~

When examining how we choose between activating or inhibiting all of these varied *desires*—vengeful, romantic, culinary, material, *ad infinatum*—there's another key question that needs to be asked: how exactly does our brain define *how much* we desire any particular thing or result (or *person*) in any particular moment? We hypothesize that there are two primary factors that combine to define our level of "desire" (aka, urge/emotional *intensity*) for any particular thing or result. The first factor is the syntactically- & vocabulary-defined overall *value* of the thing or result—determined via those three "Narrative Prioritizor Test" criteria we identified in Essay #2: importance, relevance & novelty.

Importance essentially represents that thing/result/person's experience-based *personal value* to you. In other words, if you've learned to love comic books, your brain might judge comic books as having high personal value—which increases your *desire* for lots of those particular kinds of things/results/people. Because they are based on our experiences, these importance judgements are highly individual, and *anything* can become highly important to *anyone* if their life experience has made it so.

Relevance represents the thing/result/person's "closeness to you" in terms of various in-group/out-group status judgements. If your goal impacts *any* of the social groups that you consider yourself a part of (family, school community, neighborhood, *ad infinatum*) then it has *some* level of relevance to you. So if your comic-loving self is also a member of a super-tight comic book club, or if comics are a cherished tradition that you share with your dad, then you might judge any comics-related goals/tasks that specifically impact them as *highly-relevant*. The greater the role of that social group in your life, the greater the relevance, and thus, the greater your *desire* for achieving any particular goal that benefits or relates to that social group.

Novelty represents the judgement: how many of these particular things/results/people do we *perceive to be* available vs. other things/results/people in this same category. For example, how many of these particular comic books (*Spider Man* #1) do we perceive to be available vs. other comic books in this category (*all Spider Man comic books*). The higher the novelty judgement, the greater our *desire* for that particular thing/result/person. Here again, because our *objects of desire* are heavily shaped by personal experience, the way in which we categorize those objects and how much we value their novelty is highly individual.

As noted, all of these value-defining criteria can be employed by applying a narrative parcel's syntactic & semantic content to our emotional equations. This "overall value" definition of things/results/people is the first factor in determining *how much* we desire anything in particular. The higher the overall value, the greater the desire. Thus, the comic-lover from our previous paragraphs might most-greatly desire & cherish a rare comic book received from (or possessed by) their collecting-partner dad.

The second primary factor in determining desire is something that we refer to as "emotional attenuation"—which is a cousin to what Rolls' dubs "sensory specific satiation." As Rolls describes, in the process of satisfying a general urge like hunger, we experience a satiation of sensory-specific stimuli (like a particular kind of taste) throughout the process—which drives us to diversify what we consume as we satisfy that hunger (helping us to meet different dietary needs). According to our theory, a parallel "satiation" or "attenuation" mechanism (a decrease in stimuli response intensity after increased exposure to that stimuli) is also present within our *value-based* emotions. (Such attenuation is *not* likely present within our belief-based or purely validity-based emotions—because greater exposure to belief- & pattern-prediction-supporting stimuli typically *strengthens* our subsequent response intensity, i.e., *increases* the likelihood of belief compliance, or our degree of confidence in that pattern prediction.)

There appear to be two kinds emotional attenuation mechanisms: general emotional attenuation (e.g., attenuation of general affection) and category-specific emotional attenuation (e.g., attenuation of a specific source of affection). This is, essentially, how our brain judges our "appetite" for (or our *level of deprivation* regarding) any particular thing/result/person in any particular moment.

For example, if someone has an active social life with lots of friends & family and is a frequent recipient of their abundant affection (aka, lots of exposure to that "general" emotional stimuli) the prospect of a mid-week social gathering might not

generate an urgent desire to attend (aka, a decrease in response intensity). Conversely, if someone who's starving for affection (minimal exposure to that general emotional stimuli) receives the same invite, they're more likely to have a stronger desire to attend (a more intense response to that general emotional stimuli). These are both demonstrations of general emotional attenuation mechanics at work.

Now imagine that the affection-abundant person is provided that same invite, but in this case the gathering will an include an old best friend that they haven't seen in years. Because this person's affection-response has been attenuated (by high levels of general affection) the emotional affection stimuli must be fairly significant in order to generate an intense response (aka, a stronger desire to attend the gathering). Luckily, the presence of the long-lost best friend provides category-specific affection (the affection of a dear old friend) that this person *has not* been heavily exposed to. Thus, this category-specific emotional stimuli is still capable of providing a more intense response from our affection mechanisms—leading to a greater desire to seek out that affection at the suddenly-more-exciting mid-week gathering. However, even in a case like this one, if our "general affection appetite" has been *super-satiated* we still might not be able to generate enough desire to motivate our attendance (*I'd really love to see Bob, but I'm just too burnt out from all the socializing—I really have no desire to go*).

Conversely, and sadly, if someone is deeply affection-starved, even an ultra-redundant category-specific low-value emotional stimulus—like a thrice-daily 10-second social interaction with a doorman whom you know is only really pretending to like you—can still generate a more intense emotional response, and thus, a stronger desire to ensure that the interaction always occurs thrice daily, *ad infinatum*.

The mechanics of attenuation, desire & decision-making also have significant implications in regard to both behavioral *and* chemical addictions—which are clearly major problems in modern society. In fact, addiction issues beg for an entire essay focused on them. But *Free Will* still beckons, so instead here's a quick glimpse into the unique difficulties posed by *chemical* addictions... Imagine that your stomach actually grew a little *every time* you got carried away with the joy of consumption and over-ate. It's easy to see how an infinitely-sizable stomach and its subsequently infinitely-sizable appetite could result in us consuming far more food than the rest of our systems could adequately process. This is the problem that chemical addictions produce.

Because the attenuation of a purely chemical appetite (like opiates) has no physically-correlating boundaries (like a

stomach), and because the amount of the chemical available has no production limitations (like those that physically limit the production capacity of internally-generated chemicals like endorphins)—our external chemical appetite and the amount of chemical available are nearly infinite. I say *nearly* infinite because there is one very firm boundary to our chemical appetites: death. In the end, sure, emotions like affection are powerful stuff & we can get a great buzz from large doses—but we're never going to love something so much that we generate a big enough buzz to accidentally kill ourselves. The same cannot be said about those chemicals that so many of us have grown to love so dangerously much.

One other tragic attenuation-based side-effect that relates to opiate addiction in particular: in the view of our theory, the addiction itself drastically lessens our internal capacity to exercise neurally-based *willpower* when attempting to choose *not* to satisfy that infinitely-growing appetite. As discussed in Essay #2, we theorize that our willpower capacity (which is a key factor in our ability to use narrative motivations to resist an urge) is heavily dependent on endorphins—and endorphins make use of the same kinds of opiate-based receptors that the opiate-based drugs are severely *over-attenuating*.

In other words, the same neural mechanic (attenuation) that's causing the addict to require larger & larger doses in order to achieve a response (*get high*) is also causing the brain to require larger & larger quantities of endorphins in order to engage our willpower. The major problem being that (as we noted) our brain actually has internal physical production limitations for chemicals like endorphins—so once that opiate-appetite has outsized our internal capacity for endorphin production, our willpower is utterly toothless.

In a scenario such as this, an individual's willpower capacity is essentially reduced to zero. These are likely some of the reasons why opiate addiction is so deadly, and why the widespread availability of powerful opiates in our society is pure madness. (Infuriatingly, pharmaceutical corporations' desires for increasing profits far exceeds their willingness to behave responsibly in pursuit of those profits. In particular, the Sackler family and their company Purdue Pharma have been historically vile in their greed-driven disregard for how opiates can destroy lives.) This is also why it can be crucial to provide access to "weaning" drugs like methadone—because minimally-satisfying the drug desire with smaller dosages of less-potent chemicals allows the brain to slowly *lower* that level of opiate attenuation (without the interference of intense cravings—*aka*, *withdrawal*). From our view, this helps to provide those endorphins with a fighting chance to effectively engage those willpower mechanisms in the battle against the self-destructive behavior.

Returning to our exploration of a *neurally-healthy* decision-making process, fortunately, most of our emotionally-motivated choices and behaviors are not nearly as prone to total

willpower failures—although they *are* prone to smaller ones. In fact, humans are notoriously fickle creatures whose willpower to resist or succumb (aka, *impulse control*) frequently flickers between inhibition & activation on an ongoing basis. Thus, even when a narratively-expressed intention & its selected (potential) action are initially inhibited, they can still have a powerful influence on our *subsequent* action-decisions. In other words, going back to that potential strangler, even if they inhibited their *initial* impulse to reach toward their target (which was generated via their anger-induced response "I'm gonna kill him")—they still might very well a *want* to kill their target.

This means that if, for example, the primary emotional and/or belief judgement(s) that initially inhibited that action weren't very strong (e.g., a belief that *Murder is only bad sometimes*, instead of *Murder is always bad*) then the action (or a closely-related action) might have a good chance of being *uninhibited* in a subsequent round of processing. In such a case, the subsequent parcel of dialogue that generates & accompanies the action might only be tangential (linguistically) to the original intent, and not directly intentional. For example, our aspiring strangler might simply think "Screw it!" as they now activate their just-inhibited strangling motor scripts. (Instead of thinking a more directly intentional "No, really, I'm gonna kill him.")

In our brains, "Screw it!" can also mean "I'm gonna do this!" (and carry that specific action-triggering intent) because in our narrative construction the brain has actually selected "Screw it!" *specifically to represent* the intent "I'm gonna do this!" That's because—as explained back in our discussion of *narratively contextual rule application*—the *intent* of an under-construction narrative parcel is typically shaped & directed in part by the syntactic content of the *previous* narrative parcel (when it's not an entirely new narrative thread).

Thus—even though that first impulsive urge to strangle was suppressed— *in response* to that just-experienced thought "I'm gonna kill him!" and its intent, the continuing anger (& a focus on its target) can help to direct the generation of another (& possibly more intense) unconscious impulse to overcome the inhibition and try strangling him. The brain then linguistically expresses this desire (which it has identified as a pattern in the unconscious right brain emergent data and is processing under conditions & intent partly defined by the previous narrative parcel) as "Screw it!"—because it has learned to symbolically & habitually express such an impulse in that linguistic fashion.

And if the action was initially inhibited by something flimsy (like a weak belief) then something as simple as a small increase in anger when expressing "Screw it!" could help to tip the scales in favor of enacting the previously-inhibited but now "re-triggered" potential action. Before you know it, our momentarily well-behaved *potential* strangler is suddenly (& prosecutably) an *actual* strangler.

Our brains can also tip the scales in favor of eventual action (after an initially inhibited impulse) through the *internal* use of that selected-but-suppressed motor script. As demonstrated by examples like free throw shooters who perform better after mental rehearsal, triggering a motor script while suppressing physical action allows us to internally practice—essentially, imagine—the selected action. In emotional (& thus, decision-making) terms, this practice can provide us with additional *confidence* in taking successful action (and allow us to further calibrate those motor scripts based on imagined predicted results). What are we doing when we pause before jumping across a puddle, swinging our arms, bouncing our knees & picturing our leap? What's a dog doing when it whines below a higher-up landing, springing slightly with a few mini-hops before launching itself upward? We're both building our confidence in taking successful action—internally practicing (via our premotor & somatosensory cortexes) motor scripts that have been selected, but suppressed due to judgements like anxiety or fear (or in some cases, simple otherwise-motivated intent).

Together, all of this speaks to the surprising power (& usefulness) of selected-but-suppressed actions—which can serve both as a silent instigator who recruits emotions like increasing anger to motivate subsequent action, and as a quiet coach who helps us run enough practice drills to provide the confidence to go ahead and act. (Selected-but-suppressed actions can be *so* powerful that we built our entire loop of language-based human consciousness around a selected-but-suppressed action: internally-expressed, but unspoken, speech—aka, *internal dialogue*.) The power of selected-but-suppressed motor scripts is also why—if you're really hoping *not* to lash out with an uninhibited angry action—it's probably better to focus on doing *something else* instead of focusing on *not doing* that ill-advised thing. When we focus on specifically resisting an action (by repeating a thought like "*don't lash out*" or "*don't look down*") we're actually helping that selected-but-suppressed action to stay within our *potential-action* cue—thus providing the inhibited action with more opportunities to subsequently become uninhibited.

Rather than forcing our brain's *executive* to continue inhibiting a powerful urge every time you think "*don't lash out*" or "*don't look down*," it's likely more effective to focus our mind on sustaining a preferred *alternative* action. For example, avoiding an angry action by repeating a simple alternative action-oriented thought like "*just stand here*" or better yet "*just take some deep breaths*" (which is better because physically slowing and calming ourselves can impact how we *feel*—possibly helping to dampen the strength of emotions like that *scale-tipping* anger).

In a system such as this, our DLPFC's *executive* is less of an arbiter who considers various arguments, and more like a carnival strong-man hoping to slam his hammer with

enough gusto to ring the bell at the top of the pole. Action-encouraging emotions like confidence, excitement, pride (and *anger & flight*, which cleverly employ endorphins to generate action from their inhibitory pain) —these help the strong man hammer with gusto, and action-inhibiting emotions like anxiety, fear & guilt can weaken his swing. In other words, those action-encouraging emotions can help to open & excite the neural pathways that carry action triggers forward toward our motor cortexes, and negative emotions can inhibit & dampen those neural pathways. (Rolls describes this kind of decision-making mechanism as "a single winning neuronal population.")

Although our wide array of variously-contributing (& often *opposing*) emotional judgements can make this decision-making process exceedingly-complex in humans, the process has roots in (& still resembles) some very ancient & rudimentary behavioral mechanisms. In both roundworms & humans this "action-triggering" mechanism does the same thing: it provides the capacity to override/interrupt current behaviors or actions and generate different situationally-appropriate behavior or actions when a new possible value gain/loss (i.e., a yummy or a threat) is identified.

Roundworms squiggle along according to rote, repetitive motor scripts until sensory organs identify bacteria and help generate serotonin- & dopamine-based neural responses that trigger ingestion motor-scripts, which override/interrupt the suddenly-inhibited locomotion motor scripts. Humans dance feverishly according to unconscious, musically-calibrated motor scripts until their eyes glean the delicious & thirst-quenching cocktail, and the sight helps generate urge- & narratively-based neural responses that trigger walk-to-the-table motor scripts, which override/interrupt the suddenly-inhibited dancing motor scripts. In humans, if an action trigger receives enough emotional support (or fails to generate *any* emotional discouragement) those motor instructions are carried forward via those switchboard mechanisms within our thalamus & basal ganglia (which possess connections to *task-selecting* cortical areas like that prefrontal cortex [8,9]).

All of our competing & cooperating motor script triggers—sparked by both the unconscious & internal dialogue narratives —end up sending their variety of competing/cooperating instructions through those switchboards. The switchboards then do what they've done since the beginning of vertebrates: manage & mix all those inhibit & activate messages so that the most-currently-appropriate result is eventually generated within the body's nerves & muscle fibers (via the thalamus & basal ganglia's connections to areas like the premotor & primary motor cortexes [10,11]). And tucked within this whole process there can be an additional & fundamental action-trigger that accompanies a narrative parcel: the act of

actually speaking those words aloud, instead of just hearing them in your head.

It may appear, based on the description above, that those switchboard mechanisms are actually where the final action *choices* are being made—since this is likely where all of those *varied-source* motor triggers are either *shunted-away-from* or *allowed-to-travel-to* our motor cortexes. To help clarify our theory's view of the process, I'll provide a (tragically over-simplified, but still useful) analogy of our thalamus & basal ganglia's switchboard mechanisms...

Imagine a virtuoso "train-track-switcher"—someone in charge of a massive train station with myriad tracks traveling in & out. Some of these trains carry action-triggering instructions, and each of those trains needs to have its incoming track switched over to the proper outgoing motor-area-destined track (*if available*). Most importantly to our virtuoso switcher, each incoming action-train arrives with a specific "prestige" (or *priority*) value—and the switcher's ultimate goal is to make sure that the highest-prestige action-trains get first & best access to their necessary outgoing tracks.

These trains originally depart from myriad brain areas—like our sensory cortexes, amygdala, and that decision-making *prefrontal cortex*. And although, in most typical cases, our highest-level urges can override our narrative desires—if you possess the right motivation & enough willpower, that narrative desire can still win the decision-making battle. Thus, when a high-priority action-train arrives from those PFC-based *executive control* brain areas, it ultimately seems to be capable of superseding the prestige of *any* other competing action source—as demonstrated by those *most-excruciating* fear-&-pain-overcoming and executively-ordered actions, like *cutting off your own arm to save your life* (which requires those *executive* instructions to supersede the amygdala's *reflexive* protective response).

This is how an executively-directed (*consciously contemplated*) choice can be essentially inevitable *before* it reaches our switchboards, because its instructions are typically given highest-prestige treatment within those switchboards. Any executively-directed instructions can, however, be nearly instantaneously *interrupted* & replaced by one of those reflexive protective responses if, for example, our survival is suddenly threatened by a rapidly-approaching object (an object that was not contemplated in the momentarily-previous & conscious action-decision).

There are also plenty of lower-prestige instructions that are constantly allowed to pass through our switchboards & be enacted by our motor systems because they simply don't conflict with any of the "outgoing tracks" that are required by any currently executively-controlled actions (e.g., unconsciously walking & drinking your coffee while consciously thumbing a message on your phone). In all of these

different cases, our switchboards aren't making the kinds of analytical, consciously-directed action selections that we associate with real *choices* or *decisions*—these switchboards are simply (& complicatedly) routing *previously-determined* instructions according to their *already-assigned/calculated* prestige/value and current motor-resource availability. (And these switchboards are also likely helping to route those internally-rehearsed *selected-but-suppressed* motor instructions to destinations like our premotor & somatosensory cortexes.)

In the case of lower-prestige, unconsciously-maintained actions (like drinking the coffee) it's also important to note that the data helping to direct or guide those actions is much more *generalized* than the data that helps guide consciously-maintained actions. And this *data generalization* leads to those many slightly strange or misdirected action results that litter our days. For example: You're sitting in the living room reading the paper on your iPad when you suddenly realize that you forgot to give your children their vitamins before school. *I'd better take those out & put them on the kitchen counter so I remember later.* This consciously-directed decision leads you to stand & start walking toward the kitchen, iPad still in hand…

Once you start walking and re-engage consciously with the iPad (because *of course* you keep your head buried in your tablet as you walk) those consciously-sparked vitamin-directed action tasks essentially leave our conscious loop and slip into that "action-maintenance loop" we identified earlier. And in this loop, the actions are no longer being directed by that more specific consciously-contemplated end goal (*put the vitamins on the counter*), but rather, by more generalized & immediate action-defined goals that the conscious end goal first sparked (*walk into the kitchen & open the cabinet*). Thus, these now unconscious actions are guided by that more generalized environmental/physical data regarding a path to the kitchen & a target for opening (research suggests that these kinds of "sensorimotor intentions" are specifically handled by the *posterior parietal cortex* [12]).

This is why, once you're in the kitchen and lift your head from the iPad to scan the open cabinet, you might be baffled momentarily (*what am I looking for?*) before remembering the vitamins and suddenly realizing *you're not even looking in the right cabinet*. When you first stood up from the couch, the pre-conjured task of opening the cabinet might've led you to actually open your free hand slightly in preparation for grasping the cabinet handle—leaving those unconsciously-looping instructions on continual hold until the right environmental/physical data triggers the rest of the *well-chunked* cabinet-opening motor tasks. But in our action-maintenance loop, this very *generalized* task-triggering data (e.g., the general physical attributes of a cabinet handle) is *no longer connected to* (nor does it *contain*) that more specific end-goal data

(vitamins). Thus—even though you have absolutely no doubt about which cabinet the vitamins are in—because your action-maintenance loop only needed *cabinet handle* environmental/physical data in order to let loose the rest of its *open-the-cabinet* motor scripts, you simply (& stupidly) opened the first & closest cabinet noted within your visual arena.

And—partly depending on factors like whether you cued-up that grasp & open task by opening your free hand when you stood—you might get to "*what am I looking for*?" as soon as you find yourself standing dumbfounded in the middle of the kitchen. In this case, *opening the cabinet* never even made it into your action-maintenance loop (or it wasn't part of an elaborately-chunked motor task)—leaving that loop with no further unconscious instructions to let loose once you've reached the end of your walking path. This is another example of how our mind can unconsciously maintain & direct motor actions via more generalized information *without* having any conscious awareness of how (or *if*) that action is related to our more specific (& originally-consciously-defined) end goal.

These seamlessly (although *sometimes misdirectedly*) interweaving motor instructions occurring in response to widely-varied types of data from widely-varying sources are the ultra-complex descendants of those simplistic, ultra-ancient inhibit & activate behavioral mechanisms (observed in our dapper roundworm, C. Elegans) now grown into Godzilla over the course of evolution. And another point that this all reinforces: although we have been primarily talking about the central loop of consciousness that produces internal dialogue, as a whole the many processes we've discussed in these essays involve multiple loops & offshoots (that we *haven't* discussed) that merge & diverge from that central loop. For example, after that neural moment of potential action selection, that data might also be sent to motor cortexes—priming them for potential action execution at the same time the DLPFC is determining action activation/inhibition. There is *a whole lot* happening at each step through this neural maze—we're just traversing the main perception-to-action thoroughfare. [13]

To continue bringing our essay-to-essay journey full circle, this internal neural battle between competing & cooperating instructions is reflected in a description we offered in our first essay: *Deep down in our psyche, these are the kinds of impulses that are battling for our brain's undivided attention. Each moment of existence is a Roman Colosseum in our minds—each urge, each impulse, each desire tossed into the arena, fighting viciously to be heard, to be made part of the story, to be expressed out there, where the thing that thinks them acts its act in the world.*

Revisiting The Great Syntactic Divide

Admittedly—if your hope is that yes, humans *do* possess free will—the evidence

presented so far by the *scientific* approach to this question is not encouraging. In terms of that chicken & egg conundrum, science is clearly leaning toward the side of the egg: our brain seems to have the capacity to *decide* to act (choose to activate/inhibit an action or task) slightly *before* that action's accompanying thought actually emerges in our consciousness. Returning to our act of drinking the coffee while reading, it seems clear that the coffee-sight or coffee-pang input could at least *simultaneously* spur the act of reaching *and* the internally-experienced thought "coffee." So by the time we *hear* the thought "coffee," the instruction to reach has likely already been given (or the task has at least been *chosen* to be subsequently—and essentially inevitably—enacted).

Consider that in terms of our loop, that parcel of internal dialogue's (auditory-cortex-aided) emergence within our Dynamic Core is that dialogue's *last* destination before returning to our unconscious processing. Thus, this neural moment must occur a few steps *after* that dialogue is processed by those (left-hemisphere-originating) narrative-building & analyzing mechanisms — which are necessary to first *create* & *evaluate* (syntactically & emotionally) the dialogue that we ultimately hear.

As we described, those build & analyze mechanisms are nearly *immediate* predecessors to neurally *enacting/inhibiting* that dialogue's accompanying (*possible*) action triggers via our DLPFC—because as soon as that action-triggering narrative parcel has been emotionally analyzed (for gain/loss, beliefs, etc.) the DLPFC has all of the data it needs in order to executively *enact or inhibit* those potential action triggers. Therefore, in a system such as this, our executive neural processes could receive all the necessary decision-making data *prior* to any accompanying dialogue's *actual appearance* within our Dynamic Core.

Nonetheless, there still might appear to be a tiny opportunity for free will to make a comeback in our argument. This is because thus far we've mostly explored more impulsive acts like lashing out angrily or reaching for the coffee, *not* more deliberate acts like reading the article—*aka*, circumstances where our consciousness is more focused, and therefore where it is most likely to exercise any capacity for true free will. Exploiting this tiny (& seemingly final) opportunity for free will's existence means asking the question: does the *conscious contemplation* of a task-triggering choice (via action-specific internal dialogue in a decision-making context) somehow make us *consciously aware of* that choice *before* the action is executively (& inevitably) chosen to be enacted? (Thus giving us some possible *agency* over that choice—*aka*, free will.)

This hopeful space within the loop where free will might yet exist is tucked inside the same (left-hemisphere-originating) narrative-building/analyzing mechanisms that we've been discussing. It's the location that we identified in essay #4: *The Great*

Syntactic Divide. Just in case (as it does for me) this term still only brings to mind John Belushi's delightful, garish mug (via the 1981 film "Continental Divide"—yes, large portions of my right-hemisphere are filled with peculiar data-associations) I'll briefly re-locate this epic juncture in our internal dialogue loop. The Great Syntactic Divide is that left-hemisphere neural moment *after* your parcel of narratively-based dialogue is actually constructed (via rule & vocabulary application to the highest priority pattern of emergent right-hemisphere data).

As described, this adjacent post-construction moment is when all of those emotional equations—including emotion-producing beliefs— are applied to the just-built narrative (a process that ultimately involves circuits with myriad brain areas) in order to help generate & send-off the proper instructions to the widely-varied emotion-producing portions of the brain, *and* in order to send those results to the DLPFC for use in decision-making.

We've hypothesized that these emotional equations *must* be applied here because—based on neural judgement principles supported by Daniel Kahneman's *Prospect Theory* equations[14,15]—we know that the brain's emotion-producing calculations are based on those contextual/narratively-defined relationships between value & validity. Therefore, the brain *can't* apply its emotional equations until *after* this data has been constructed as a narrative. Beliefs are also specifically-designed to make judgements about narratively-constructed data, and would therefore naturally need to be applied *after* this Great Syntactic Divide as well.

Where Are You, Free Will?

What opportunity does this provide for true free will's existence? Not much. If we genuinely have *agency* over our choices (dialogue-based *conscious awareness* of choosing to act prior to the executive—and inevitable—neural *activation* of that action) then internal dialogue would need to *circumvent* this Great Syntactic Divide, sending special & specific *action-impacting* dialogue immediately into our Dynamic Core to be heard *before* the brain undertakes all of that decision-determining emotional analysis. This would seem to be the only way to become *aware* of an inevitable choice to act *before* our executive machinery receives the emotional data that determines whether any action is activated (or set into *inevitable* motion). And this does not seem to be a very plausible data pathway for our internal dialogue.

In fact, our theory has already provided some very strong evidence that internal dialogue *must* be emotionally analyzed before it is actually heard by us or spoken aloud. This is because all dialogue—even sentences that we only hear in our heads—includes some type of *inflection*. Angry, perplexed, sad, curious, excited, timid, etc., etc., etc.—no matter what the words are saying, their *inflection* almost always expresses some type of *emotion* when they

are heard by us or spoken aloud. And, as discussed in the previous essay, in order to properly calibrate the inflection of dialogue according to emotion (something that likely involves Brocha's speech area) the dialogue obviously needs to be emotionally analyzed first. This means that our brain's executive decision-making mechanisms (which make *immediate* use of this emotional analysis) are likely set into inevitable motion slightly before that *(inflected)* dialogue can actually be heard by us.

Nonetheless, although our last ditch effort already looks dead in the water—in a final nod to free will's tenacious elusiveness—we'll explore one more full example of how consciously-experienced internal dialogue might *appear* to precede & activate an (inevitable) action. *(And we'll see if it really looks anything like the kind of free will you were hoping for.)* We can navigate this neural territory by considering a scenario: one of those *dodgy* married men is being heavily wooed by an attractive "other woman."

Observation of the woman's overtly & explicitly flirtatious behavior serves as raw data input that emerges from the right hemisphere as a pattern from which the left hemisphere constructs the internal dialogue: *"This chick wants me."* In the micro-moment before he hears himself say this—just after the Great Syntactic Divide—this man's emotional equations & beliefs might generate a combination of excitement & guilt that he *feels* as he has this thought. Has he decided if he's going to kiss her yet? Probably not quite. After this hypothetical observation is processed, the *next* thought might generate (among other neural items) the act of possibly kissing her. The enacting of this "kiss her" motor script might be inhibited by a number of factors beyond just the possible guilt, or the fear that his hypothesis is wrong. Matters of consent, appropriateness of setting, insecurity over one's kissing skills, a sudden craving for a cheeseburger, etc.—all can be factored in via sophisticated syntactically-based rule & vocabulary application, belief application, urge analysis, and emotional equations.

Let's say our potential couple is alone in a private room, and after she impatiently queries "*Are you going to kiss me or not?*" our married man finally *does* choose to kiss her. Did he choose *via free will* to violate his beliefs and kiss this clearly-non-wife-of-his? Well, let's say he *said to himself* "Okay—I'm just kissing her" in the moment *before* leaning in and planting one on her lips. And yet, if he was actually deciding to act *as* he was saying that to himself, he would (by the definition of action being ultimate proof of an *inevitable* decision) *already* be kissing her (*while* he was thinking this) and not just *talking to himself* about it (yet not actually *acting*). In fact, in the moment after having that thought and *before* actually kissing her he could *change his mind* and resist the impulse despite previously declaring internally his intention to act.

Thus—even if he consciously & linguistically expresses his "I'm just kissing her" *intent* in the moment *before* kissing her (and not *during*) —the moment in which he actually neurally *sets into motion* the "kiss her" motor script (his activation of that now-inevitable action) likely occurs just after the construction of the *subsequent* dialogue parcel. This means that his actual *choice* to kiss her still happens (as it *always does*) at that point just after all the necessary narrative-evaluating, decision-impacting data has been calculated & made available for use in choosing to activate or inhibit an action: within that DLPFC *executive* area that immediately follows all of those post-Great Syntactic Divide emotional & predictive calculations—and a couple steps before the accompanying dialogue arrives (fully-inflected) in our conscious awareness.

The conclusion seems inescapable: the neural *instruction* (or choice) to act is triggered via mechanisms that are likely enacted *just microscopically prior* to any consciously-experienced *awareness* of that choice is capable of being expressed or observed *via internal dialogue's* subsequent auditory-cortex-aided emergence in our Dynamic Core's multi-faceted arena.

Even if our married man skipped *all* of that internal dialogue foreplay & kissed her *immediately* in response to his first perception of her *wantingness* —the resultingly-triggered neural *instruction* to act would still depart from the same loop location & still occur *just microscopically prior* to hearing that perception's & action's "spurring" or *accompanying* thought. In reality, the *actual thought* that accompanies a moment of true action & decision like this is often less about *intention* than it is about *experience*: "This is awesome" or maybe simply "Oh my"—or probably something more along the lines of *"This is a bad idea."*

Damn You, Science

Even after exploring the hinterlands of our Great Syntactic Divide, it appears that science's verdict remains unchanged: we *decide* to act in the micro-moment *before* we *think* about the act (or have whatever thoughts that *appear to us* to happen in conjunction with, or somehow spark, choosing that action). This conclusion—drawn here from closely examining the looping neural mechanisms that we've proposed—is also supported by Benjamin Libet's famous research, which has detected within individuals apparent decision-making neural activity that precedes the subject's conscious awareness of choosing the action.[16]

Even though our brain truly does make choices *based upon* how we perceive ourselves to feel about the situation & our specific individual beliefs, our *awareness* of our brain's calculation of those factors in our decision-making process occurs slightly *after* our brain actually makes those calculations —calculations that are, in reality, what determines whether an action truly becomes *inevitable*. To which my ultimate response is, for the most part: *so what?* And I say this because we have yet to finally answer the *first*

question we posed about free will, the philosophical one: *what truly is free will?*

In the end, when we humans express our desire to "have free will," we are essentially saying that we want to believe there is a true "Agent of Self" within us, an agent who *is* us, and through whom we consciously—and without any undetectable internal interference—control the choices, actions & behaviors that our body physically enacts. And I believe that our primary objection (or even *repulsion*) to the idea that decisions precede our conscious awareness of them is rooted in the belief that such a system robs us of that true Agent of Self. In other words, if the person that I *perceive myself to be* arrives onstage after the script has already been written, then my *perceived self* is merely an actor, and not a self-directed agent of any kind.

But what really *is* free will? Although there is no single-entity, fully-perceivable Agent of Self (it's actually our *whole system*—conscious & unconscious) we do, indeed, act as we feel we want to act—as long as we are *capable* of acting. Is this not essentially the core claim of free will? Does it matter whether or not we can consciously perceive *why* we are sometimes incapable of enacting our will? In some cases, a deeply desired or intended action is inhibited in ways that we are entirely conscious of (like a marathon runner who desperately intends to, but simply cannot, take that next step) and sometimes a deeply desired action might be inhibited in ways that we are not entirely conscious of (like desperately wanting to kiss someone who desperately wants you to kiss them, but simply not being able to act). In both cases we are still *aware* of our conflicting desires & capacity, and we feel that those factors ultimately reflect our inner *agency*—regardless of whether we actually can or do enact our self-expressed intention.

These mechanisms create a system that, *as a whole*, behaves exactly as any fully-autonomous Agent of Self (*with a sometimes limited capacity to enact its self-expressed will*) would behave. I like to refer to this conundrum—that we feel & behave exactly as if we had a genuinely autonomous, Agent-of-Self-driven free will, even though that solely-conscious, all-controlling Agent doesn't really exist—as the *Free Will Paradox*. This paradox is probably best expressed simply by the conclusion: *humans have, for all intents & purposes, genuine free will, except that they technically don't.*

Determinism-Schmerminism & The Truth About "Morality"

This brings us to one last philosophical notion that, like qualia, has held the spotlight for far too long in the arena of brain theory: *determinism*. Basically, determinism suggests that in the absence of something like a true "Agent of Self," all of the decisions & actions made by any creature with a purely physically-based system of mind would be *pre-determinable*—if you just

happened to know *exactly* all the conditions that will contribute to this decision.

In other words, according to "determinism" *any* purely physically-based system of mind is ultimately a totally robotic & fully *agent-less* creature (since our choices are ultimately "merely" the result of at-least-momentarily "pre-determinable" neural & physical responses to our cornucopia of data input—as if that *miracle of existence* was something worth lamenting). Of course, this whole question of determinism has one fatal flaw: it has nothing to do with *reality*. In reality, there is absolutely no way that anyone could know *exactly* all the conditions that will contribute to a decision—which in many cases are virtually *uncountable* when you consider the myriad synaptic connections that are hit or just-missed in every round of thalamocortical processing (in every sensory, internal & cognitive system) and that contribute to every thought, action & *interaction* (which also exponentially increases result variables) that occur over a lifetime (which has been spent accumulating millions of unique data modules that also impact the predictability of those decisions).

The extraordinary neural complexity, malleability, "inter-causality" & "re-programmability" of this cognitive process are also central to the brilliant Peter Ulric Tse's anti-deterministic argument for free will in his book *The Neural Basis of Free Will: Criterial Causation* [17]. Ultimately, the entire *idea* of determinism is so uselessly irrelevant to anything that might relate to actually understanding or shaping human behavior that applying its principles to our *actual existence* is nothing less than totally absurd. The way that over-thinking (yet oddly short-sighted) philosophy-types try to make the question of determinism relevant to *actual human existence* is by saying crazy things like "if determinism governs all behavior, then we have no moral justification for punishing criminals, since their actions were not truly *chosen*, but merely the inevitable result of who they are & the situations they encountered." But this kind of logic is all twisted up in something else that philosophers talk a lot about, but that doesn't *really* exist: *morality*.

What the crazy statement above is trying to point out is that if determinism *does* truly rule the mind then punishing criminals is essentially *immoral* because they are not really *at fault* for their actions—therefore the criminal's actions themselves cannot be considered actually *immoral* (which is dubiously circular logic to begin with, since it means we also had *no choice* but to imprison them, making the whole question of whether we should or shouldn't moot). But *punishing* criminals for the innate *immortality* of their actions is *not* really why we imprison them. Essentially, we imprison them because—based on their actions—we predict that they are likely to act in this socially-destructive way in the future. Although throughout our cultures we talk about criminal systems being *punitive*—and

founded upon an idea that criminal acts *deserve* to be punished—in the end, this is just our cultural & personally linguistic way of contemplating & expressing those deeply innate neural impulses to make decisions about protecting ourselves from future losses by assessing known evidence & applying reliable predictive patterns. (And as we mentioned in our previous essay, the notion that we use prisons for *reforming* offenders is simply a rather obvious *lie* that we pretend to believe because it makes us feel better about ourselves.)

In other words, despite what we say, we don't really imprison someone *just because they murdered someone & murder is immoral.* If we did, then there would be no such thing as *innocence due to self-defense* (or *innocence due to wearing some sort of official uniform while killing people*)—which we've conveniently declared as "moral" acts of murder. The real purpose, however, of such exceptions (to our brains) is to help sort out acts of murder that *are not* supposedly good predictors of future socially-destructive behavior, and thus *do not* require punishment to achieve a socially-desirable result.

Deep amongst our symbolic, contextual neural calculations—in the same way that a gain or loss is not about the *money* or *object*, but rather, how a change in our access to its *perceived* value will *help or harm* us—criminal punishment decisions are not based on judging the *morality* of the *act* of violence or theft, they're based on judging what those acts *predict* about *future behavior* that might help or harm us. And when certain societies or communities tend to punish, for example, specific races more harshly & frequently than the general population—even when committing *the exact same acts*—it's evidence that those decision-makers' minds are biased toward *predicting* that the "demonized" race is more likely to cause future societal harm. (Any individual's communally-nurtured brain-logic varies according to whatever data they've consumed.)

We've naturally built our system of societal rules & responses exactly how our brain works: by basing those rules & responses around making the best *predictions* about the *most desirable* or *most undesirable* future results (aka, the most valid & beneficial gain & loss predictions). From this perspective—as opposed to being a cause for *doubting* the reasoning behind imprisoning criminals—determinism is *exactly* why we should imprison people like violent criminals: because the predictability of behavior means that removing these individuals from general society is highly-likely to lead to an ultimately desirable result for that society, which is the primary goal of sequestering people like violence-prone individuals from the rest of us. (We also try to make everyone in society *aware* of potential punishments because we *predict* that the *fear* of such consequences will likely *prevent* at least some potential offenders from committing socially-harmful acts in the first place, aka, *deterrence*.)

Of course, because of those uncountable aspects of every decision, a *truly* determinism-based "Minority Report" & cognitively-predictive (somehow *neurally*-based or *gene-predisposition*-based) justice system is an entirely impossible *fantasy*. This means that we have to wait until someone *actually* does something awful to make a good prediction about whether or not they are likely to do something awful again in the future, and thus decide if they should be sequestered from the rest of us *apparently-much-less-likely-to-be-awful* humans.

Note that we didn't have to mention *morality* anywhere in *our* explanation of crime & punishment. Because there *is* no static or timeless *morality*. There are only culturally-developed, individually-learned *beliefs* that guide how we ultimately judge the "moral content" (aka, *social benefit*) of any act. And beliefs are not about any inherent *morality*—they're just a very special version of that thing our brain is obsessed with: a prediction. Beliefs are merely high-validity, high-value prediction tropes that help to guide our actions toward an ultimately desirable (gain-enhancing/loss-averting) result.

Morality, *per se*, is simply a non-starter where the brain is concerned—it just doesn't correlate to how our cognitive systems manage & judge data. All of those "moral" behaviors like aiding & sharing, affection, empathy, *not cheating* (on a test *or* your spouse), *forgiveness*, *etc.*, *etc.*, *etc.*—every apparently *self-sacrificing* or *purely-other-entity-benefitting* human action has been accounted for in our theory by some emotion, belief system or other *contextually-framed*, *survival-supporting*, *data-based & evolutionarily-arrived-at* neural mechanism. Of course, the ultimate societal *result* of applying all of these systems in a communal fashion over many millennia is exactly the same thing as what we consider to be a "moral code." Although—because the nature of our neural systems means that such "moral codes" are deeply & broadly culturally-based—the notion of morality is actually the *opposite* of what it's typically considered to be: in truth, *morality is highly malleable.*

Bring Me My Soul

In the end—even after dismissing determinism—Narrative Complexity's own kind of *paradoxical free will* is still not enough for most of us; we simply *want to be* that singular, fully self-informed & self-determining Agent. And as we stand on the deck of this once-wayward *Free Will* vessel now finally in the harbor, hollering into its empty hold for everything that it *has not*, at last, brought home to us—what we are really saying is very simple: we want a *soul*.

But what would a soul really be? Wouldn't the behavior & choices of a creature with a "soul" be the same as we behave & live right now? Can't this extraordinary, exquisitely-evolved & unimaginably complex system of mind be *equivalent* to a soul? Isn't each of our *minds* something that is uniquely us? Something purely based on a mix of unique

inborn attributes, unique experiences, feelings, thoughts & desires, and uniquely acquired & organized rules, beliefs & vocabulary? Is this not, for all intents & purposes, exactly what a soul purports to be? Almost. But it does, of course, lack something very fundamental that is what we really seek from a soul: permanence. The human mind *cannot* give us permanence.

Forever does not appear to be something that these magnificent & genuinely *soulful* neural mechanisms & systems can give you. And, frankly, that pisses me off. I *want* some kind of permanence. I want *to see what happens* and *be part of it*. I *am not* "okay" with my non-permanence just because I won't actually be able to *perceive for myself* the horror of my non-consciousness. I am utterly terrified by a state of being that I will never know.

And so, for me, this is what I have: my *desire* to be here—which seems like an ultimately irrational (or at least overly-circularly-logical) motivation for being. *Be to be.* And yet I *believe in* and *cling to* that desire to be here. For me, this fervent, life-defining & ultimately-irrational desire is the closest thing *that I have* to a soul. It *is* me. The thing that says *I am* and *I want to be*. The thing that someday will be *was*—a thing that I wish would be capable, in that someday, of saying: *"I was. I was."*

Direly, based upon everything I have so far learned in this life, and all of those self-defining beliefs, and rules, and words, and experiences—I simply do not believe it is likely that this mind will ever someday say to itself, *"I was."* We live, I believe, by definition, in the universe of the present—it is the only place we ever truly are, or will be. We exist now.

And I believe that in all the ways that might genuinely matter—we, the unique being & mind that is each of us, *do* have free will & full domain over the choices that we *perceive* to be perpetually presented to us. It is simply that a great deal of that unique being & mind is making its contributions to our ongoing self *behind the curtain*. But the work that goes on unseen is just as much a part of *who we are*—that unique amalgam of self-accumulated & self-organized data—as who we consciously *perceive* ourselves to be.

I am here. We exist now. You are a mind in the present. The most extraordinary expression of *self* that this earth has ever created. You will know nothing but this, but you may try to know as much of this as you might desire while you are here. Within the confines of your circumstances, you may choose to do with being here whatever you wish—even to quixotically battle those confines, to seek to alter the world in which you roam. That freedom, this mind, its temporariness, and the will to do, and be—these are what we have been given.

The will of the free, and a mind for the now. A place we are in time. The melancholy glory of being.

###

Endnotes:

p. 204:
1. Bargh, John A., and Ezequiel Morsella. "The unconscious mind." *Perspectives on psychological science* 3.1 (2008): 73-79.

2. Morsella, Ezequiel, et al. "Passive frame theory: A new synthesis." (2016).

3. Edelman, Gerald M., and Giulio Tononi. A universe of consciousness: How matter becomes imagination. Basic books, 2000.

p. 205:
4. Deacon, Terrence. The Symbolic Species. Norton, 1999.

p. 209:
5. Wymbs, Nicholas F., et al. "Differential recruitment of the sensorimotor putamen and frontoparietal cortex during motor chunking in humans." *Neuron* 74.5 (2012): 936-946.

p. 214:
6. Domenech, Philippe, and Etienne Koechlin. "Executive control and decision-making in the prefrontal cortex." *Current Opinion in Behavioral Sciences* 1 (2015): 101-106.

p. 216:
7. Rolls, Edmund T. *Emotion and decision making explained*. Oxford University Press, 2014.

p. 223:
8. Klein, Johannes C., et al. "Topography of connections between human prefrontal cortex and mediodorsal thalamus studied with diffusion tractography." *Neuroimage* 51.2 (2010): 555-564.

9. Alexander, Garrett E., Mahlon R. DeLong, and Peter L. Strick. "Parallel organization of functionally segregated circuits linking basal ganglia and cortex." *Annual review of neuroscience* 9.1 (1986): 357-381.

10. Kultas–Ilinsky, Kristy, Elena Sivan–Loukianova, and Igor A. Ilinsky. "Reevaluation of the primary motor cortex connections with the thalamus in primates." *Journal of Comparative Neurology* 457.2 (2003): 133-158.

11. Marsden, C. D. "What do the basal ganglia tell premotor cortical areas." *Motor areas of the cerebral cortex* (1987): 282-300.

p. 225:
12. Andersen, Richard A., and Christopher A. Buneo. "Intentional maps in posterior parietal cortex." *Annual review of neuroscience* 25.1 (2002): 189-220.

p. 226:
13. Haller, Matar, et al. "Persistent neuronal activity in human prefrontal cortex links perception and action." *Nature human behaviour* 2.1 (2018): 80.

p. 228:
14. Kahneman, Daniel, and Amos Tversky. "Prospect theory: An analysis of decision under risk." *Econometrica: Journal of the Econometric Society* (1979): 263-291.

15. Tversky, Amos, and Daniel Kahneman. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and uncertainty* 5.4 (1992): 297-323.

p.230:
16. Libet, Benjamin, et al. "Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential) the unconscious initiation of a freely voluntary act." *Brain* 106.3 (1983): 623-642.

p.232
17. Tse, Peter Ulric. The Neural Basis of Free Will: Criterial Causation. MIT Press, 2013.