An Infinite Game

by Matt Illingworth

Abstract

In order for humanity to avoid existential threat - that is, to prolong our death as a species indefinitely - we must develop a new approach to the threat of existentialism as it pertains to life. In this paper, I posit that our difficulties as a species have their root in a basic misapprehension of the infinite; that our ambition to grow forever in size, ability, and accomplishment reveals a powerful discomfort with the concept of infinity, which compels us to play finite games for an ostensibly infinite length of time. Since we seem not to have any intention of ending play in said games, prudence suggests we either alter the rules of engagement or invent new games entirely. I outline various possible methods for restructuring our scientific, socioeconomic, and philosophical games in an attempt to reconcile them with the infinite game we call the human endeavor.

As finite entities living finite lives, humans naturally find themselves at odds with the concept of infinity. Though we may often feel a sentimental affinity for the infinite as imagined in the form of a loving and ineffable God, the infinite - as both an abstract concept relevant to the human experience, and as a ubiquitous and oft-maligned result in math and physics - stubbornly continues to vex our understanding.

The very appearance of infinities in our equations and our collective imagination demands attention and explanation, but out of our inability to make sense of an ultimately senseless problem, we design mental and societal constructs that either deliberately ignore it or bargain with it: In the religious case, we quell the panic of mortality by assuming that a finite number of actions taken in life equate to the characteristics of an infinite afterlife; in the scientific case, we demand that infinities ideally have no suitable role in the equations that describe our universe; and in the existential case, we attempt the despairingly futile task of filling an apparently infinite void of inherent meaning with a finite amount of secular meaning.

To paraphrase Aristotle: The mind, like nature, abhors a vacuum. Hence, our existential dilemma looks like a universe of human intent feverishly expanding to fill the eternal emptiness surrounding it - and the concept of even decreasing the speed at which we transcend becomes absolutely repellent. There exists no inherent reason for us to grow our scientific knowledge, our capacity to eliminate suffering with technology, or even our markets and physical presence in the universe - but since we appear fully intent on doing all of these things regardless, the task of humanity becomes one of changing our organizational schemes - both sociologically and

technologically - to allow for levels of complexity that tend toward infinity.

This means that the difficulty inherent in social engineering lies not just in the sociopolitical biases of the designer; as blogger and complexity theorist Nick Szabo explicates, in order for societies to grow in complexity - which includes all the varieties of growth mentioned above - some kind of economy must exist to facilitate the expression of said growth, for the simple truism that proliferation of knowledge and abilities necessitates role specialization. It follows that those who would reject currently-prevalent economic models on purely moral grounds fail to account for how the replacement model would handle the complexification attendant to increased prosperity.

Social scientists, then, charge themselves with the dual task of maximizing prosperity and civil liberty on the community level, while maintaining the organizational integrity of large-scale economies as prosperity and complexity increase to infinity. Because such a model would necessarily straddle the line between sociology and data science, and because there exists no obvious way to predict how the aforementioned complexity will manifest, I humbly suggest a few closely-related areas of study that may prove useful: the co-evolution of human brains and tribes, and the computational complexity thereof.

Until fairly recently, neuroscientists imagined that the human brain operates like a very densely-layered computer, with all of the deterministic reductionism - comparing neurons to circuits and lobes to computer chips - that such an idea entails. In a popular interview, philosopher and cognitive scientist Dan Dennett discusses his and others' shift from this ill-conceived notion to something far more dynamic and strange: That the brain behaves like a complex network of separate cognitive functions which vie for attention and influence, form alliances and make mutually-beneficial agreements, and so on.

If we accept this premise, then we find a natural isomorphism to the behavior of governments, economies, and the internet - and hence, we can gather all manner of dynamic change within these systems into a more general way to think about thinking: The memes that virally propagate through a society - fundamentalism, reformism, totalitarianism, etc. - find their physiological correlates in the neural interactions that give rise to various aspects of consciousness - resistance to change, sudden enlightenment, and aggressive narcissism, to name a few. However, if we believe that each neuron behaves like an individual while simultaneously acting in concert with a hundred billion of its peers, we must ask why only 7.5 billion humans can behave much the same way and yet perform so poorly overall. To put this another way: If there exists an impairment in the collective consciousness' cognitive processes, from what

underlying condition does it arise, and how do we attempt to remedy it?

The implicit assumption in any answer to this question has to do with the shape of societies themselves; when communists, socialists, and free-market capitalists argue as to how social programs should function, each posits something akin to a data structure or computer network, whose shape determines the qualitative experience of human life. The methods by which parts of a society communicate with each other mirror the algorithms which might operate on these structures - and the necessity of said algorithms exposes a problem that continues to go only partially rectified: We did not evolve to live in such large groups.

In the early 90s, anthropologist Robin Dunbar put forth the idea that the mean size of a particular primate species' neocortex correlates to the size of that primate's social groupings. According to Dunbar, the ideal population of a human tribe rests at close to 150 people, where each member knows and interacts with every other member and social grooming takes a significant portion of group time[1]. Though popular science writers often take Dunbar's number as the maximum amount of people the human brain can keep track of, it seems a more fruitful approach to consider this number as the maximum social complexity the human mind can model[2].

This means that social groups may operate at their highest efficiency only when each member's brain contains an image of the over-arching hive-mind - more explicitly, communities greatly benefit from each member's understanding of every sub-community within the whole[3]. The apparently holographic nature of the mind - where the brain manages to retain functionality with as much as 90% of its mass removed - would seem to affirm this assumption on a different scale; the way in which brains store distributed copies of themselves hints at an implicit knowledge of the connections between each neuron, like a network graph stored in each node of that same network graph. One can then extrapolate this to the expression of group dynamics, and analyze how communication breaks down in large populations.

It appears that the flaw endemic to growth in human population lies in our inability to socialize and coordinate with groups larger than Dunbar's number; in terms of the holographic metaphor above, this equates to the brain's inefficacy at modeling the relationships that comprise arbitrarily large group-minds. The solutions posed by modern political systems commonly use hierarchy as their organizational blueprint - but the enactment of these structures often leads to decoherence between the respective priorities of representatives and the represented, because the complex web of relationships between colleagues exhausts high-level politicians' capacity to model group dynamics. Capitalist societies ostensibly mitigate this error by establishing markets, whereby the underrepresented can influence large communities with their money - but doing so

requires a degree of financial security unheard of for most people, which further guarantees the supremacy of hierarchical institutions[4].

Even worse, in lieu of some social programming that would otherwise handle various statistical outliers in the phenomenology of tribal community - violence, theft, land disputes, etc. - we institute various systems of justice and prosecution, which tend to reinforce the hegemonic biases of hierarchy through sheer violence. In all of these cases, the social software that replaces communal self-organization falls prey to corruption, decay, and egomania. This last consequence highlights an underlying conflict of differing human timescales: the person motivated entirely by ego attempts to play finite games where infinite games would be more appropriate.

According to philosopher James Carse, participants in a finite game assume the role of "player" and attempt to end the game by winning. Players of finite games agree on a set of rules by which they decide the victor, who then carries a higher status in subsequent games. In contrast to this, Carse discusses the additional concept of the infinite game; where participants play by altering the rules indefinitely, with the goal of avoiding resolutions and continuing play forever. In Carse's model, all aspects of life fall under one of these two categories. One would include chess, politics, and even whole societies in the former category, with pursuits such as culture, art, and hackey-sack in the latter. Finite games very often come at the expense and detriment of infinite games: the unscrupulous, despotic politician oppresses and imposes unjust restrictions, hindering the expression of cultural advancement for personal gain; the ruthless, amoral billionaire accrues an ever-inflating fortune, claiming "social darwinism" while limiting the economic evolution that the poor could otherwise contribute - thus homogenizing the capitalistic gene-pool and weakening markets' chances of survival.

Interestingly, capitalism as practiced today straddles the line between these two classes of game; as a roadmap for the future of humanity, capitalism depends on infinite growth - but the most powerful and successful capitalists tend to play finite games out of pure self-interest, which diametrically opposes infinite growth. Moreover, any strategy of revisionist, egalitarian capitalism would still have to contend with the finite amount of terrestrial resources at our disposal - and beyond that, with the apparent finitude of the observable universe.

One could argue various morally-driven political points about this quandary, but it seems more interesting to examine potential solutions from a perspective of pragmatic technicality: If the complexity of human activity must increase with prosperity, and we assume that the conventional paradigm of tree-shaped hierarchy eventually fails, then the density of connections between neurons in the human brain may offer an alternative model to strive towards. This

would entail several things: a dismissal of monolithic corporatism and monocultural values, and the adoption of grass-roots communal organization; rejection of traditional value-based politics in favor of practical, goal-oriented cooperation; and the union of local and global community through the space-binding effect of the internet. Social policy decisions would thereby propagate from the bottom-up, based on the priorities dictated and mediated by interpersonal behavior - and macroscopic social coordination would no longer befuddle our innate capacity to store community relationships in conceptual models. I conjecture that a system like this, with its carefully-chosen resemblances to Dan Dennet's "community-of-mind" model, would provide the durability necessitated by continual complexification of society.

One might imagine various ways of mapping this new web of human organization to weighted graphs, based on the exchange of goods, services, information, communal time, or even seemingly trivial things, like frequency and subjective quality of idle chit-chat. Clusters and superclusters of densely-connected nodes could then represent the individuals that constitute local, national, and digital communities; the sum-total of which could then also contain itself in the form of many copies distributed throughout the network[5]. One might balk at the unprecedented loss of privacy such a system would call for; but it would appear, given the ever-advancing state of cyber-warfare and big-data technology, that true privacy may soon become a thing of the past.

Fascinatingly, the bitcoin protocol appears to satisfy the transparency of such a societal-network model while maintaining individual anonymity. To obviate the necessity of trusted third-party mediators in financial transactions, peer-to-peer exchanges between bitcoin addresses become a matter of public record, which the network then verifies by consensus - but the identity of each user remains obscure. One could argue as to the ultimate longevity of bitcoin specifically, but the democratizing technology of crypto-currency seems to herald a future of human organization similar to the one imagined in this essay. One might also object to the bitcoin protocol by pointing out the possible "51% attack", in which economic terrorists or felonious mining pools take over more than half of the network; doing so would grant any attacker the ability to either destroy the market or spend money they don't have. This issue merits special consideration if one hopes to build a durable future on these technologies, and it naturally leads to broader questions about the cybersecurity of technological utopia.

If we chose to build a digital hive-mind based on crypto-currency technology, on which the entire human social operating system depends, the importance of information security could far outpace its own capabilities. Further, if neuroscience progressed to match our science-

fictional fantasies of cybernetic telepathy and literal group-minds, the integrity of our very identities and perceptions would become perilously dubious. It follows that the future of humanity relies heavily on the protection of the hacker community - a body of people with a strikingly liberal attitude towards moral relativism - each of whom openly admits to assuming the role of player on (at least) one side of a finite game. Can one reasonably expect the beneficent coders to prevail unerringly over the malicious? To answer one question with another: Can one reasonably expect the black side of a revolving yin-yang to progress faster than the white? The two opposites move in tandem; no final victory or defeat can occur.

Thus, the 51% attack reveals its congruence to the rogue neurons that sew discord in unhealthy minds, or the haphazard gambits perpetrated by finite players in the high-stakes game of government and finance. The myth of utopia presumes the infallible scrupulousness of each citizen, but students of history learn to expect an element of irreducible rascality in all human activity[6]. In the language of Carse, resolutions never take place in an infinite game -which makes impossible any attempt to end the harmful play of finite games. To reiterate an earlier point: There appears not to exist any obvious, inherent reason to improve the species or ensure our infinite longevity - except for desire, for the pure human will to transcend. But since we appear fanatically eager to expand our technological powers indefinitely, we obligate ourselves to accept all the inexorable negative ramifications. Going further, beyond the realm of ethical quandaries, the problem generalizes to a truism about dynamic evolution of societies: as members of a species whose complexity tends toward infinity, no social programming can remain effective forever.

In a discussion about humanity's stewardship of its future, it seems sagacious to point out that any proposition offered - including the ones presented here - has an expiration date, a limit beyond which it either loses usefulness or degrades in performance as a result of corruption. The cybernetic hive-mind outlined above might very well handle an increase of complexity for a time - if we managed to deal elegantly with ever-advancing cyber-threats - but this too would soon fall short of the infinite growth we desire so strongly. One cannot do away with the limitations that infinite ambition automatically imposes, and the realization of this invokes the existentially enervating feeling of climbing an infinitely slippery hill.

If the prospect of eternal existence and unlimited power (for humanity as a whole, if not for the individual) instills a sense of world-weary fatigue in the reader, consider the "Technological Singularity", a hypothetical expression of the infinite in human history that nevertheless has a boundary in time. This concept implies that the sequence of events that drive

scientific discovery grows asymptotically with time, culminating in the birth of something wholly unpredictable by the human mind - but there follows from this argument a question that remains unanswered: What exactly would such an acceleration look like in a relativistic universe, on a finite planet?

Einstein tells us that an object approaching the speed of light does so asymptotically, while simultaneously approaching infinite mass. If the complex dynamical systems that compose scientific progress could be digitized by an artificially intelligent supercomputer, and then somehow accelerated to near light-speed, this machine might well collapse into a black hole and thus obfuscate the information contained inside - if this acceleration even proved achievable in the first place. This should be taken as nothing more than speculation, but physicists have good reason to believe that the interaction of singularities with phenomenal reality, when it occurs at all, tends to have catastrophic repercussions - recall the reality-shredding effects studied in black hole physics, or the negation of physics at the moment of the Big Bang.

One could make a similar argument about the finitude of matter and energy, on earth and beyond; the explosive growth of human population requires us to one day leave the planet, because finite resources cannot support infinite growth - and even if humanity managed to colonize other galaxies, the material universe would inevitably obstruct the boundlessness of human initiative. The appearance of asymptotic lines towards a finite point in the future should instill fear in the mathematically-savvy - for even the mere suggestion of singularities in human development implies calamity, a fatal acquiescence of the finite to the infinite.

Mathematicians of various disciplines refer to functions that tend to infinity as “divergent”, and one can perform various tricks to obtain convergent results from divergent expressions; meaning that infinities need not break our means of rational understanding if we handle them with care. Whatever strategy we ultimately adopt for navigating the future, we would do well to incorporate a similar renormalization of the infinite sum called “the future of humanity”; undertaking such an adjustment would entail attenuating the current rate of economic and technological growth, and cultivating a heightened willingness to play infinite games. It may strike the reader as an unattractively mundane solution, too far removed from the flashy, aspirational immediacy of start-up culture and transhumanism - but it appears absolutely essential to concede our limitations as a young and naive species with inordinate power. To state it differently: Perhaps we need not grow up so fast.

Concomitantly bolstered and maddened by what lies just beyond the horizon, humanity’s insecurity surrounding the boundless future catalyzes a futile compulsion to anticipate and profit

from it; but to play an infinite game requires effortless submission to unremitting change, not the rigid proceduralism and duplicity more natural to finite games. Consequently, the human issue of comporting ourselves with integrity has a direct correlation to infinity: our societies psychologically prime us to crave resolution, but no such event ever occurs in the ceaseless process of becoming. This inner disharmony engenders in us a death wish, a yearning for escape from the trap of infinity - which, if acted upon, ultimately represents not a victory over the infinite, but a submission to it. The only permanent solution lies in the whole-hearted acceptance of transience; in coming awake to the structural hinderances in our collective psychology; in abandoning those finite games that squander our potential and our creativity; in learning to play - and take pleasure in - the infinite game of life.

1 To give the reader a quantitative idea of this: If every person has some connection to every other person in a population of 150, a simple binomial coefficient calculation $\binom{150}{2}$ gives 11,175 as the number of relationships in the group. In a population of 220 – another popular estimate of Dunbar's number - the same calculation gives 24,090.

2 Social network theorists have shown some evidence for this idea by looking at the average number of teammates per group in various MMOs (Massively Multiplayer Online games). The reader can learn more about this here and here.

3 In terms of graph theory and computational complexity, this equates to an upper bound on the brain's ability to perform the clique cover problem on a graph of the tribe. While such a calculation falls under the category of NP-hard, the human brain's capability of interpreting social cues makes the problem much simpler than the brute-force approach a computer would take.

4 A data scientist might refer to these social models as some form of weight-balanced tree, with power and/or money represented numerically and concentrated mostly at the top. One might find some method of analyzing the economic health of a society by looking at the number and shape of the branches between the root and the leaves.

5 Biologists and mathematicians might do well to study closely the dynamics and organization of bee-hives; bees' extraordinary talent for difficult mental math and the efficiency of their community efforts may have some special significance – a hidden link that could provide the blueprint for the human-hive proposed in this essay.

6 To illustrate this point with the tools of set theory: One could imagine civil rights movements as the consequence of excluding oppressed communities from the set of all people with rights and privileges. However, with each new addition of groups to this set, more groups emerge that find themselves excluded from it – a truism made all the more poignant in the coming age of biotechnology.

Reference

Szabo, Nick . "The Irreducible Complexity of Society." *Unenumerated*. N.p., n.d. Web. . http://unenumerated.blogspot.com/2006/02/irreducible-complexity-of-society.html

Dennet, Daniel. "The Normal Well-Tempered Mind." *THE NORMAL WELL-TEMPERED MIND*. N.p., n.d. Web. . http://edge.org/conversation/normal-well-tempered-mind

Dunbar, R.I.M. (June 1992). "Neocortex size as a constraint on group size in primates". *Journal of Human Evolution* **22** (6): 469–493. http://www.sciencedirect.com/science/article/pii/004724849290081J

Forsdyke, D. R. (2009). Samuel Butler and human long term memory: Is the cupboard bare? Journal of Theoretical Biology, 258, 156+

Carse, James P.. *Finite and infinite games*. New York: Free Press, 1986. Print.

Nakamoto, Satoshi. "Bitcoin: A Peer-to-Peer Electronic Cash System." *bitcoin.org*. N.p., n.d. Web. . https://bitcoin.org/bitcoin.pdf