

UCLA

Lake Arrowhead Conference, 2005

Title

On a Hierarchically Decomposed Agent

Permalink

<https://escholarship.org/uc/item/0bh3m6vp>

Author

Suarez, E. Dante

Publication Date

2005-09-13

ON A HIERARCHICALLY DECOMPOSED AGENT

E. Dante Suarez
esuarez@trinity.edu

Trinity University
Department of Business Administration
One Trinity Place
San Antonio, TX 78212
Phone: 1 (210) 275 5916
Fax: 1 (210) 999 8134

1.- INTRODUCTION: DECOMPOSITIONS AND LEVELS.	1
2.- UPPER and LOWER LEVELS.	5
3.- US vs. THEM.	14
4.- OPTIMALITY IN A HIERARCHICAL WORLD.	19

ABSTRACT

The present manuscript is intended to informally elucidate my ideas on a general theory of collective behavior and structure formation, with a resulting architecture that can be broadly applied. The proposed model represents a decomposition of intent, based on the idea that an agent's behavior, whether it represents an individual or a group, can be seen as an emergent property of a collection of intertwined aims and constraints. I consider a disentangled agent that is formed by multiple and relatively independent components. Part of the resulting agent's task is to present alternatives, or 'fields of action' to its component selves. Correspondingly, the composed agent is itself constrained by a field of action that the superstructure to which it belongs presents. The superstructure of agents possesses a certain amount of *cohesion*, and can thus be ascribed agency and modeled as a unit; its independent parts could be consciously or evolutionally constructed and aligned.

This hierarchical representation intends to reveal behavior as bounded and bounding, exemplified by an agent that is abstractly defined. The most important aspect of such an agent is that it maximizes its objective function by definition, given the constraints of the nature of its components and the superstructure to which it belongs. According to this proposition, suboptimal behavior is often misclassified because we do not recognize the actual maximizing active agent.

The flexibility of this model can be used to describe groups or societies in interaction. It also allows us to understand experimental evidence that humans cooperate more than they should according to theories that define them as unitary. The theory to be developed would explain complex organization taking into account all aspects of subset reorganization, which depend on the particularities of the entity. This view must include a measure of group welfare beyond Pareto optimality, as well as definitions of the concepts of a *behavioral function*, *optimal scales and design*, *bordered maximization*, *cohesion*, and *identification processes*.

Keywords: Hierarchical decomposition; Agent-based modeling; Pareto optimality; Utility and Welfare functions.

1.- INTRODUCTION: DECOMPOSITIONS AND LEVELS.

This essay poses what I consider to be salient questions about the connections between the approaches to describing human behavior and social phenomena offered by economics, social theory, and evolutionary biology. To describe a more general theory of social interaction, I propose a hybrid model in which I drop customary assumptions made in economics about what is considered a decision-making unit. In this view I recognize that a person is composed of many different intertemporal individuals, and that groups may behave as a unit. More explicitly, a person's incentives now may be in contradiction with the same person's incentives a moment later; alternatively, a military coalition, racial subgroup, or social class may behave as a whole with properties that emerge in the aggregate.

Traditional social science culminating in the Enlightenment concentrated on the study of entities that were clearly delineated, where one could separate what belongs to a person's nature against the backdrop of what does not. By the same token, the agent is taken to be exogenous, and therefore disconnected from the agglomerate to which it belongs. I take a poststructuralist view in which we move beyond what is referred to as 'the myth of a unitary subject.' Standard economic representations have focused on individuals as the smallest unit or the ultimate irreducible atom of the paradigm, but I propose further attention to the possibility that such units may actually be agglomerates, the products of internal networks that deserve attention. For example, we can think of a company's organization or a strategic military coalition, and no agency is allowed for identifiable players forming the organizations and alliances, nor emphasis is given to the study of blurry abstract borders dividing the participating members. The possibility of understanding some of the internal complexities that give rise to their positions is essentially ignored.

One of the main purposes of this research is then to bring to the attention of the field of economic modeling the need to break apart the threads that make up an agent, whether a person or company, both in an intertemporal and a static sense, in order to recognize the emergent properties that compose it. Most importantly, if we accept this proposition, then some of the founding theorems in economics must be reformulated, particularly those based on Pareto optimality. An investigation into a hierarchical decomposition reveals that important margins are overlooked when the physical person is the unit of measurement in defining a Pareto optimum and its corresponding underlying efficiency. The arbitrariness of the use of the static and isolated individual as a unit of measure can be demonstrated by asking ourselves simple questions: Shall we refer to ourselves as the one we are in this very moment? Should we stop caring about our children?

The research thus attempts to redefine what a unit of decision is, as I claim that one could unscramble behavioral influences to the point of not being able to clearly delineate what the individual is, who is part of a group and who is not, or where a realm of influence and action ends; the boundary between an individual self and its social coordinates is dissolved. The agent can be thought of as a person, a family, a social class, a political party, a country at war, a species as a whole, or a simple member of a species trying to survive. By the same token, if we are to redefine collections of individuals as operational units—such as ants forming an indivisible anthill—then we must rethink what we mean by ubiquitous concepts such as altruism and selfishness.

Throughout this work, I describe different instances of abstract concepts of 'hierarchically defined levels', and to draw analogies it is imperative that I use welcoming terminology. The word most commonly used to describe an agglomeration is 'group'. However, that might bring to mind a set of humans, and the concept must at all times be kept at an abstract level. 'Set' is also not appropriate because it forces the differentiation between those that are members from those that are not. This is because individuals or groups of individuals may wholly or partially belong to an agent, in many different coordinates and according to drastically different definitions. The word 'network' is a good candidate, since it brings to mind a system interconnected with varying intensities. Instead of the commonly used concept of 'boxes', the word that comes to my mind is 'bubble', since I picture them to be sometimes joined at the border, one inside of the other, or having an intersection. The proposed blurry agent will be defined by a host of definitions, and will therefore include some components according to one criteria but not another. In this sense, the analysis of the North Atlantic relations in the twenty-first century proposed by figure 1 will have a large number of levels.

Without an obvious winner, I use several of these words to convey the sense of an agglomeration that does not have clearly defined boundaries. All in all, these ‘entities’ must be seen with a broad eye that recognizes not only the elements of physical connectedness, but also those that may be abstract, but just as causal. In terms of social groups, for example, I will think of them as both subjects and objects, in the way we would think of the chemical field of an ant that is both controlling the individual ant’s behavior, but also is determined by the ant’s feed. There is much to draw, however, from the word ‘boxes’, as it is the preferred term in hierarchical descriptions, where a group is inside a superstructure, like an endless chain of smaller boxes inside bigger ones. As I present such a hierarchical characterization, I would like to keep such visualization. However, I must stress the fact that the clusters I describe may have ‘porous’ borders, and just as a fractal structure may not allow us to separate what belongs inside the box from what does not.

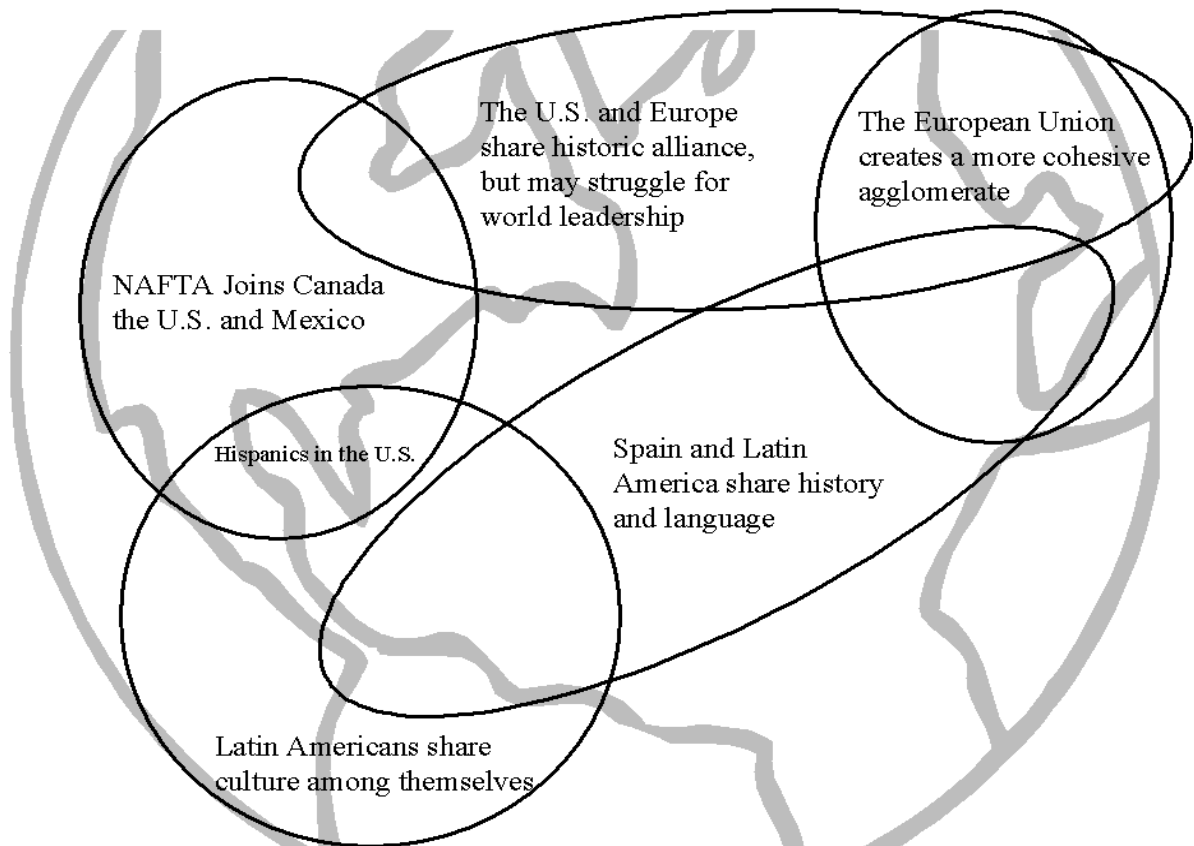


Figure 1

In characterizing where a Mexican-American agent would lie in a simulation, we must consider that it would be classified by more than one category. Characterizing the players of a realistic strategic interaction will likely leave us with overlapping definitions. In this figure we can see that no group has a border clearly delineating it from its counterpart in all relevant coordinates, just as the United States now has a growing Hispanic community that could be both defined as part of Latin America as well as part of the U.S.

A crucial discussion in social theory has revolved around the degree of agency allowed for conceived subjects. The discussion is about the definition of subjects and objects—and the corresponding separation between interior and exterior—about whether we consider the individual as the subject or the origin for all social relations, or favor a model that understands the individual as an object that is a product of such relations. Although my position is that we cannot truncate a subject from its circumstances, from the object in which it is defined, I claim that the distinction is a quantitative, not a qualitative one, so that we shall be able to measure some degree of agency. Who is the active agent in a stock market model in which the buyers and sellers have very low levels of cognition, but are kept in check by the topology of a well-developed market? In this case I claim that the upper level that the market represents is an agent that can adapt to circumstances, perhaps by design, perhaps because of how it evolved, or perhaps even because of the intrinsic order it embodies as the epiphenomenon of the selfish but traders with low levels of cognition.

In this fashion, the agents as they are commonly defined, such as people, may not be thought of as necessarily possessing much agency, as they represent machine-like structures that react to the incentives presented. Nonetheless, even if we could conceive of a universal function describing every single decision humans make, (and engage in a discussion similar to that of Albert Einstein's assertion that God does not play dice with the world,) the deeper understanding of concepts such as chaotic and deterministic, sensitive to initial conditions, and the general unpredictability of complex phenomena now allow us to conceive of a subject-object relationship in which the combination of both concepts does not make up for the aggregate they compose. Therefore, I take in this work an intermediate position of double causality between the parts and the whole, as has been proposed in complex systems. Or as the French psychoanalyst Jaques Lacan puts it, there can be no separation between the subject and society, for society inhabits each individual, and we can only draw from the cultural repertoire available to us.

As we attempt to address these and other related questions, it is imperative that we have a clear understanding of what each one of the abstract levels described actually mean, for one can otherwise fall into a trap of ascribing qualities that do not actually belong to them. A level must not be a simple reflection of the interaction between lower levels. In other words, an aggregate may seem to be ruled by an external and perhaps conscious force, while in reality the patterns it describes can be completely attributed to its participant components. In this sense, we must distinguish between an upper level, and a simple environment in which the agent finds itself. In sociology, the individual is ascribed little agency when compared to the group or social structure; classical economics, on the other hand, grants zero agency to upper level creatures, as the selfish actions of individuals are carried by an invisible hand to a plateau of overall organization. As it applies to evolutionary biology, I find this distinction to represent the core of the controversy between individual selection theory and group selection theory. Even in the absence of any agency, an abstract upper level may be subject to evolutionary pressures, and therefore potentially malleable by its subcomponents. This essay envisions a framework in which one could describe the different dimensions in which behavior is determined and in which evolution occurs, including some that can be traced back to the gene level, some that affect the individual, and some that affect a group of individuals. If we are to consider internally-disconnected people as well as some groups as agents, then the question becomes: to what degree are the decisions taken by such an agent? To what degree is an upper-level agent simply an epiphenomenon of the actions of its subcomponents, and to what degree is it established enough to make its subcomponents behave according to its abstract objectives?

The model proposed would then subdivide observed behavior into many different levels of optimization. At each one of these levels, the actor chooses its best alternative given the set of possibilities it encounters, that is, its phase space. The phase space will include transformations made by an upper level to which it belongs. In this fashion, we can think of a rowdy married man as having a utility function of his own, but also belonging to an upper level relationship that restricts his raucous nature. Moreover, the agent is composed of lower level subcomponents that may possess agendas of their own. It is the agent's responsibility to present its subcomponents with phase spaces with individual optimal solutions that are agreeable to the upper level encompassing agent. In other words, the subcomponent agents will optimize the phase spaces in which they find themselves, while the upper level must consider the manipulations of these realms of possibilities that will render the desired aggregate behavior. In this sense, if we consider a firm to be an agent, then this level is composed of the subdivisions that form the company, which in turn are run by a group of people. The company also finds itself in an upper level that includes all possible regulations for the industry, which is in turn encapsulated in a given society.

If we think of a couple as a higher level for two loving adults, the implication of the proposed paradigm is that the couple, as an entity *per se*, will have purposes which can be thought of as relatively detached from both participants forming the couple, for their individual utility functions are intertwined in the complexities of sharing, and can be returned unrecognizable from their origins. Moreover, if the couple is 'sanctified' by a social, political, or religious institution such as marriage, then we can more clearly see that at some point both parties could want divorce, but it may be best for them to stay together because of social pressures, and it is in this sense that we may consider the couple as independent of its forming parts, and having what we could model as 'desires' of its own.

To build the proposed framework, we begin with basic individuals that are composed of sub particles that may be modeled as deciding how to organize themselves ‘upwards’ (if only at a local level), but that are only composed of parts that have no deliberate decisions at an individual level. Aside from this ‘lowest level’ in which the units have no decision game, all other levels display agents that are composed of lower level (and possibly conscious) components that are constrained by the realm of action the ‘upper level’ allows. The composing elements of any ‘lower level’ could always conceivably consider themselves as relatively independent and proprietors of a utility function; in the other direction, the level’s agents are also determined by the ‘upper level’ to which they belong, whether they recognize it or not. I do not assume agency at any of these levels, but I also do not assume that it does not exist. To make this distinction clear, I will not rule out the possibility that our kidneys work because they feel happiest doing so, nor that a government can be just as conscious as a person in deciding what is in its best interest.

The simplifications of the past were generally adopted for purposes of tractability, and understandably so, since models grow exponentially more difficult to understand when we consider the interactions of agents defined in different dimensions. With the advent of complexity theory, however, we can now imagine the possibility of tackling these problems directly, with the use of neural networks, cellular automata, numerical approaches, and a whole lot of computer power. The complexities of the proposed architecture can be endless. This notwithstanding, the paradigm for a new pandemic and inter-disciplinary science built in a hierarchical structure would accept the intercommunion by means of a parsimonious model that is broad enough to accept the fractal nature of realistic agents, but at the same time tractable enough for the capabilities of an appropriate agent-based simulation. The model intends to advance the development of a common language in which novel ideas can be transmitted across disciplines. Such a language would allow us to compare a model in which disentangled humans in a given culture have some degree of independent agency, but are also to some degree objects of their social circumstances—to another one in which countries are trying to position themselves in the evolving global arena, but are nonetheless fighting with their internal political differences, as well as with established international norms and existing treaties.

The task ahead is a long one, but I claim that the necessary tools are already in place. We are headed to the creation of a general theory of hierarchical agents. The intention of this manuscript is thus to provide an outline of what the new paradigm could look like, and I divide its presentation into six related sections. The following section (2) proposes a different approach to the understanding of an individual, in which I suggest a class of utility functions that would place the individual’s decisions in the broader social network to which she belongs, and consider the possibility that she can be subdivided in a way that her actions can be understood as the emergent property of a complex agglomeration of desires. Section 3 proposes a theory for group formation, discussing the elements necessary to understand the separation of out-group antagonism and in-group caring, and the corresponding representation that a generalized hierarchical description of society could embody. Finally, the concluding section 4 brings together the topics covered to investigate for possible definitions of optimality in a hierarchically decomposed world, and propose a general new perspective for economic science; one that analyses social bodies in light of the proposed paradigm of complex divisibility, reconsiders the usefulness of Pareto optimality when the divisions of the social network are not clear, and reflects on necessary considerations in defining possible measurements of welfare for a social agglomeration.

2.- UPPER and LOWER LEVELS.

Using the language commonly used in economic science, I explore in this section the nature of individuality and collectivism, of how individuals organize to create evermore complicated organic groups and societal structures, and how so called units can be decomposed into conflicting subsets. My intention is thus to redefine what a decision-making unit is in order to disentangle behavioral influences. However, in my new definition, I recognize that we will generally not be able to clearly delineate what the individual is, or where a realm of influence and action ends. The proposed model redefines agents in two ways. First, there are no obvious atomic agents, for all actors represent the emerging force resulting from the organization of (perhaps competing) subsets. The subcomponents in turn form an internal system that is actively reorganized, and shall be referred to as the ‘lower level’ of a structure. On the other hand, agents

are to be described within a group to which they belong, which will be defined as the ‘upper level’ of the hierarchical representation, and will constrain its subcomponents’ behavior.

Microeconomics, or Economics applied to the description of the behavior of an individual agent, has usually relied on simplistic definitions of what represents an indivisible actor in the two main levels of interaction: Individuals or households for consumption decisions, and firms for production ones. Generally speaking, firms have a straightforward behavioral directive: to maximize their profits. Correspondingly, individuals attempt to maximize their utility function, which is for all practical purposes an ordinal mathematical representation of their happiness. Textbook neoclassical economics has thus modeled behavior through a closed objective function that the agent attempts to maximize, subject to the constraints imposed by the exogenous environment. The origins of such a utility function, as well as its realm of action, are the focus of my intended research, as I find the standard Walrasian description of the utility function in isolation—disconnected from its origin, culture and society, as well as its relation to the utility functions of peers—incomplete. I consider a more descriptive utility function—one which is the result of a collection of relatively independent aims—where the agent’s ‘aggregate’ decision level represents an emergent phenomenon.

The proposed insight is that, aside from finding a proper objective function for an abstract agent, its behavior will be determined by a constrained ‘field’ that the conglomerate to which it belongs proposes. The conglomerate or upper level will need to actively arrange its subcomponents so that they find individual maximizations that are in accordance with its overall objective function. We thus find sovereign entities inside the agent that need to be controlled, and who do not necessarily maximize the objective function they would in more isolated circumstances. The analysis will distinguish between a standard objective function describing a simplified, unitary, and clearly-defined agent, and a function that actually describes how the agent reacts to any given situation, once we realize that the true actor belongs to an upper level, and that it has only limited control over itself. In line with the current expansion of the Behavioral Economics field, I shall refer to this broader objective as a *behavioral function*. It implicitly explains what the agent does under any set of circumstances, in a decision-making process that may most often consider the effect of her actions on future selves, offspring, family, peers, and all sorts of lower- and upper-level layers of which she is composed and to which she belongs.

In this sense, when we arbitrarily zoom in and analyze a relatively well-defined agent, we may classify its behavior as suboptimal in relationship to its own abstract objective function, but only because we would be artificially studying it in isolation, or without regard for the struggles of its internal nature. The upper level may force lower level members into behaviors that are only optimal for the former, and it is in this sense that we can understand the behavior of a soldier. The ‘agglomerate individual’ will have to organize its subcomponents in order to maximize its objectives, since suboptimal internal coalitions can materialize, in the way a drug addict cannot overcome the desires of a strong drug-craving self, or a tyrant may force a society to extinction.

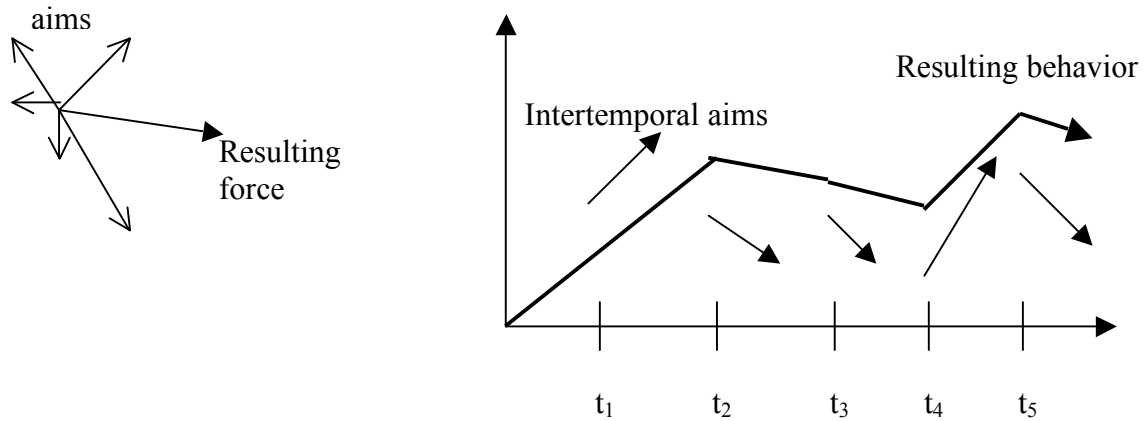
Moreover, if we redefine collections of individuals as operational units, then we must redefine what we mean by ubiquitous concepts such as altruism and selfishness. I propose a benchmark position in which all actions are revealed as selfish once one understands what the benefited acting agent is. Traditionally, we began with a clearly defined agent and tried to understand its actions as a maximization of objectives given constraints. In the proposed paradigm, we assume maximization occurs, and then work towards the delineation of the benefited entity involved. We could use this rationale to express the idea that social structures benefit a subset of society, in a conceptualization that could be coined as revealed group or social preference.

2.1 LOWER LEVELS: THE DECOMPOSITION OF A UNIT.

In this section I lay the foundations for what an individual could look like in a view that does not assume that so-called wholes cannot reorganize, or that its possible reorganization would not have describable consequences. As in standard economics, I refer to individuals by a utility function, which is a mathematical construct that orders different consumption alternatives according to their desirability, but

only in an ordinal fashion. I want to consider a more general type of descriptions of human behavior that consider the unavoidable aggregation the person represents, as well as its undividable nature as part of a social group. Figure 2 presents a graphic description of the kind of utility functions I will be considering, once I relax the stringent assumptions of invariability and inability to decompose.

Such an understanding of an individual as an agglomeration of relatively independent aims can serve as a starting point for a broader theory that describes the formation of complex hierarchical objects, stressing the in-group vs. out-group dichotomy as it applies to economics, evolutionary biology, sociology, social psychology and political science. Just as the utility function represents the desires of an agglomerate, the groups which build upon networks of 'lower level' collections of individuals may be thought of and modeled as having 'a mind of their own' as well as relative control over the individuals in its realm of action, as they act according to predictable patterns related to the maximization of their possible expansion and survival probabilities.



A. The static decomposition of wants

B. The intertemporal decomposition of utility

Figure 2

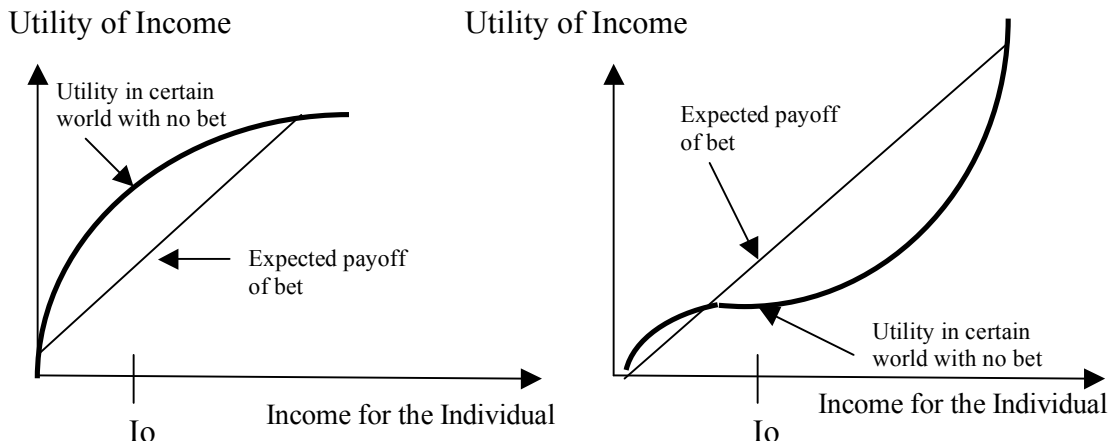
I relax standard assumptions about the utility function and allow for one that can vary though time and is the product of a complex internal process representing the wants of the disconnected subcomponents of a person.

Many people would feel uncomfortable with an introspective view that reflects a multiplicity of selves, and would rather hold the idea of a true, core self. As I describe in further detail, I conceive of the proposed legion of selves as naturally related, and perhaps having a more static and institutional complex function ordering their preferences; that is, we can think of individuals as possessing an internal government body that may be relatively stable. A different but related setup would have more disjointed selves that have recourse with related similar selves, which are perhaps called upon by similar circumstances. In terms of having different 'selves' inside, most people I know cite the one that they are in the very moment they talk about this subject, and claim that the current spokesperson is the only version of themselves that exists; granted, sometimes this person can be angry or on a particularly good mood, but the variation does not grant them in any way as having multiple personalities. The view I propose recognizes that a person plays very diverse roles in life, and that each role can be represented as an independent agent that embodies different 'subject positions.' A subject position incorporates both a conceptual repertoire and a location for persons within the structure of rights for those that use that repertoire. Once having taken up a particular position as one's own, a person inevitably sees the world from the vantage point of that position and in terms of the particular images, metaphors, storylines and concepts that are made relevant within the particular discursive practice in which they are positioned.

Significant research has been devoted to the analysis and expression of tastes that change over time, where individuals can be though of as temporally disconnected, and may face intertemporal strategic situations in which the person can play games with future selves. Much less research of which I am aware has been devoted in economics to understanding the utility function as the result of an internal (and perhaps

complex) process involving subsets of the individual entrenched in a momentary negotiation. Instead, the standard utility function is modeled under the unrealistic assumption that it is static and representing the preferences of an individual who may not be subject to relevant subdivisions. An enhanced utility function could rank all alternatives against some form of internal set of budget constraints, and provide the option that maximizes it. However, this internal process is traditionally modeled as linear and static, thus making it implicitly irrelevant. In actuality, the process itself may be highly complex, rendering the study of an individual with conflicting desires incomplete and thus inaccurate.

For example, consider the fact that many folk buy a disproportionate amount of lottery tickets, and altogether engage in many risky activities that do not fit the standard description of a risk averse individual. The people in question drastically enjoy buying a \$2 lottery ticket with an expected payoff of less than fifty cents, but would never in a lifetime consider a fair bet for a million dollars. A standard representation of this behavior can be embodied in a utility function of income that has decreasing returns to scale in most regions, but highly increasing returns to scale for larger income levels. This set up would ensure that the individual is not prone to take even a fair bet that will lose or earn him ten thousand dollars, but would be willing to do it if it implied possibly getting a few million. Figure 3 depicts such a tradeoff.



A. For normal amounts the preference is for a certain amount of income over the expected result of the bet.

B. For the possibility of an extraordinary amount of income, the individual prefers the expected payoff of the lottery bet over the certain income of not buying such a ticket.

Figure 3

Here, I depict the standard conception of a utility function that could account for the way in which individuals incur in highly unfair 'lottery' bets, but would not accept significant fair bets that could hurt them. The corresponding utility function is marginally decreasing for 'normal' amounts of income, but marginally increasing for extraordinary large amounts. Because of this particular construction, the utility of the expected payoff for flipping a coin for a thousand dollars is significantly beneath the utility of not having the bet. The expected utility of buying a lottery ticket, however, is above the utility of having the certain (perhaps low) income one has without it.

The necessary calibration of the particular utility function necessary to obtain these results is not negligible. In other words, it is not easy to conceive of a utility function that would be optimized by the behavior at hand. Moreover, such a person would have an odd view of life, seemingly reaching satisfaction at some level of income (where the marginal utility flattens out), and then suddenly finding himself in a frenzy for more money, as increases in income gradually increase his utility by larger and larger amounts. Most importantly, it is not only the implausibility of the standard utility's interpretation of lottery consumption what makes it inappropriate, it is the fact that it tells us nothing about the internal process that may give rise to such seemingly incongruous behavior.

Another interpretation could be solely based in the way that the network of disconnected selves is organized, in a setup where all selves that are not positively affected by an eventual winning ticket are so close to indifference that do not veto buying the low-priced tickets, while serious lobbying efforts are

devoted by the self with the millionaire dreams. Each one of the internal actors possesses a utility function that is nonlinear, with a threshold that makes them not care much about petty purchases. As a whole, most humans act this way, taking a long time to decide upon a big purchase, but not caring much about whether or not to buy bubble gum. Once in a while, the lottery-crazed sub agent comes along and proposes to the person's internal parliament whether to buy in a sweepstakes. All of the other agents then make a comparison of the world with or without the two dollars the ticket costs, and see no general difference in their possible well-being. Therefore, no one objects to the plan. The original proponent of the bet however, will be very happy with the purchase, and will start planning about how he will gladly spend the first prize winnings, as well as all the praises he will get from peers, just for existing. This agent will be happy with a ticket in hand, and no potential agent will have anything to say about it. Moreover, the happiness that he creates from himself can be so great that may have surplus, and conceivably affect the person as a whole. Quite contrary to this situation is the one in which a considerable sum is at play. Under such circumstances, the potentially affected sub selves come out in force to reject the demand for a bet.

If a model with endogenous utility functions is to be conceived, the way in which the internal decision-making process takes place must be analyzed. Only when we understand this internal design will we be able to understand how people make their choices, how they think of themselves, and how they join to form groups. Moreover, the processes that I would like to address in this essay go beyond the way in which the human mind is structured. From now on, I will conceive of any possible agent—whether an animal, a family, a university, or a beehive—as an emergent phenomenon possessing a relevant internal system, and I shall refer to it as a lower level.

In many cases, the lower level may seem irrelevant because it may not be readily apparent. As an upper level agent coordinates its subcomponents, its lower level may seemingly disappear, leaving us once again with unitary subjects. In such cases, the lower level components enjoy no agency, just as the soldiers of an army would if they were perfectly trained to follow all orders. In other words, our livers might have a utility function to maximize (making them happiest when they correctly process the body's substances), but that objective function has been perfectly coordinated with upper level objectives. In this sense, it is not relevant to think of the way in which human cells work when thinking about that person's behavior. However, I will envision cases in which the lower levels may become exceedingly relevant as they have the ability to reorganize, and may even actively reorganize depending on circumstances. What makes the attention to the lower levels interesting is not only the possibility but the frequency of the tendency to reorganize. Therefore, the agents I am considering are not time or situation invariant.

2.2 UPPER LEVELS: AN INDIVIDUAL AS PART OF THE SUPERORGANISM.

The previous section described the lower level of any agent composed of parts that have some degree of agency. The proposed perspective implies that decisions cannot generally be modeled as the result of an individual's fixed utility function because they are actually the result of a complex competition between decisions that are made at a lower, or deeper, level, thus putting the definition of the individual—the decision-making unit—itsself into play. Once again, the idea behind a model of hierarchical agents stems from a view of the world in which emergence is pervasive, in which we find wholes that are irreducible to their parts, and therefore exist in different dimensions. In such a world, the independence of events that is assumed in classical statistics theory no longer applies, and a holistic approach is necessary. To complement the story, I need to discuss how the agents I am considering form part of upper-level super agents. In such a view, I think of a bee as having lower-level and internal sub agents that compose it as an organism, at the same time realizing that the bee itself is incomplete if not defined as a part of an acting beehive, just as a neuron cannot be fully described outside its complex brain network. In this sense, a beehive reacts to changes in the environment, with the individual bees acting more like neurons rather than decision-making agents in their own right.

Moreover, to fully describe behavior in such a general model of hierarchical agents, we must also understand the encompassing situation facing the decision maker. In this way, we may reveal what is special about the optimizations that agents are carrying out and that ultimately define the agent itself. In other words, think of the forming of a company from a group of people previously working separately on a

given product. Even though the constructed company may have some agency problems, it has a board of directors whose orders are closely followed. The point I would like to make in this section is that this company that will be designed according to somewhat strict guidelines, spanning from the way in which the different departments are organized, to the accounting rules it will be judged by, to the regulations it will abide by, all the way to the competition it must be made prepared to face. So long as all these restrictions are reflective of an economic or survival process that increases the probability of success, I shall refer to all of these constraints encircling the company, or any other agent, as upper levels.

In terms of the utility functions that I have been considering, my approach now leads us to define the institution of a person. The idea is that the subagents, aims, or desires that compose the more general and unified perspective in which we normally conceive of humans as not entrenched in a lawless competition for attention, but rather as confined within a sort of political system, full of laws, traditions, history, habits, coalitions, and all imaginable aspects of a full-fledged bureaucracy of mind. This internal system may be well established and generally permanent, perhaps reflecting what we think of when we say 'I'. From the point of view of the lower level, the person's system represents an upper level that offers a rugged landscape for the subagents to maximize. In this sense, the desire to go on vacation is always there, but in order to be actualized, it must form a coalition based on the current political circumstances, as well as wait for the opportunity when the protocols allow for a vote on the proposed bill. Figure 4 depicts the utility function of individual i , internally composed of a vast amount of subagents with their own and relatively independent agenda. The process of defining a person is not complete if done in isolation, so the idea is that the individual will conceive of herself within a longer time horizon of intertemporal selves, as part of a family, a society, and so on. These upper levels present a rugged landscape that allows the agent to position herself where she wants, in a process that makes some regions more attractive than others.

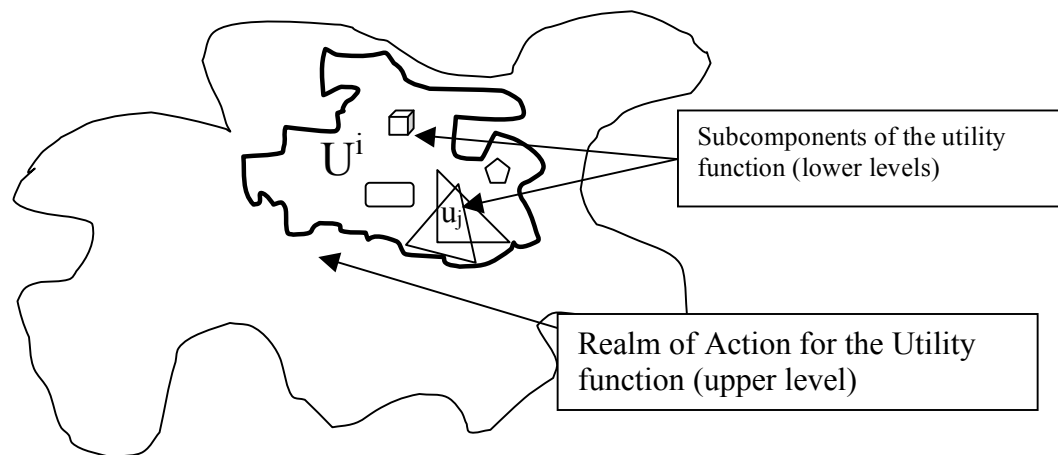


Figure 4

In this figure, I conceive of a human composed of many desires, u_j , that together with the person's established operating system compose her utility function, U^i . Such a function is in turn not defined in a vacuum, but instead fills a niche within the superstructure to which it belongs. The utility function is only an abstract description of the agent's nature in isolation, or under some definition of 'ideal conditions.' In reality, agents are defined within the system to which they belong. Such an upper level presents price vectors corresponding to each possible role that a lower-level decision-maker may choose. In this sense, a human may have a propensity to become a criminal, but society will make it harder for him to be so.

The upper levels may be created by a conscious decision or by an evolutionary force. In the case of the company, there exist organizational schemes that align the incentives of the administrators with those of the stockholders. Therefore, a conscious attempt to mitigate agency problems will create the optimal structure for the upper level: the corporation. On the other hand, upper levels may emerge simply as a product of continued interaction and a corresponding increase in cooperation, as was the case for World War I soldiers long stuck in opposite trenches who did not shoot to kill, but rather attempted to wound each other lightly, therefore providing the enemy with a much-desired ticket home. This last example represents an upper level in the sense that the observed behavior cannot be explained in isolation. Moreover, once such a level

appears, it creates a self-sustaining recursive system, which under the appropriate circumstances can be robust.

An upper level is to some degree an adaptive unit with an intrinsic design that cannot be reduced to its agents, but rather to the design of how those agents are organized. On the opposite end of this spectrum, the individualistic paradigm assumes that group organization is a reflection of self interest. The question then becomes: whose self interest? How small is the indivisible agent? Upper levels can therefore represent coordination, identification with others, institutions, implicit laws, religion, or a credit bureau. An environment, however, is not by itself an upper level, for the environment may be independent of the creatures that inhabit it. An upper level is the superstructure that creates the individual, and without which the nature of the individual would have to be redefined. In this last sense, the environment is an upper level to the degree that it is interconnected with the lower-level agents it breeds. Under this definition, both levels at play coexist, are complicit, and represent indispensable components of the recursive process that gives them life. To warrant the need for a separate discussion of these upper levels in terms of behavior, all that we need to establish is that the current “self” of an agent is not an autonomous decision-making unit—it is a subagent that must negotiate with other subagents in the committee that determines the behavior of the agent. Any actor is therefore constructed and defined by the rules of the game that its predecessors instituted, and it is also established in a reality that is constrained by what the games need in order to perpetuate themselves for future generations.

I will attempt to further clarify the way in which upper levels appear in the remaining sections of this manuscript. For the time being, consider for example a possible trend showing that regions with distinct seasons (up to a reasonable degree) tend to develop societies that are richer than those established in tropical climates. Humans living in tropical environments generally have access to invariant food sources year round, and therefore can live with relatively separated intertemporal utility functions. On the other hand, a human living in an environment with changing seasons must adapt to the fact that food in winter is too scarce for him to survive without savings. Therefore, if the summer version of the self does not prepare for the cold days, the individual will be dead before year’s end. For such a person, an intertemporal upper level must be created for reigning in the loose and selfish intertemporal selves. Such a broader conception of the self is the basis for thrift, work ethics, and ultimately for economic growth. If we are to consider the utility function of the individual working on a beautiful summer day, we must then not conceive of it in isolation, but as part of the limits set by the intertemporal society of selves.

2.3 BEHAVIOR IN A HIERARCHICAL WORLD.

Microeconomic science has been built on the pillar of Revealed Preference, which basically states that the only information we can draw from are the preferences of individuals. The concept is an important one, since it reminds the researcher not to ascribe spurious qualities to the individuals studied and abstracted by a utility function; such a function is a mere representation and may contain aspects that do not actually reflect the individual’s nature. Aside from the ongoing discussion this work elicits about how to define an individual, (and whether or not this would be a transitory entity), I find this practice essential in any model of human behavior. In other words, our only data is the actual preferences observed in the real world. In that case, how do we explain a world full of love and hatred in a paradigm of selfish individuals? Why is it that all major religions call for loving at least a subset of all creatures? How can we explain the fact that currently all developed countries have a progressive tax system?

I propose my conception of hierarchical agents as a means for—if not explaining—at least classifying and describing the way in which “us vs. them” lines are established. In my model, the agent can represent a level of strategic decision making, or the evolution of adaptability and responsiveness. The agent confines the subcomponents that belong to it, but at the same time it is constrained by the upper levels to which it belongs. I believe that an agent-based model constructed with the architecture outlined here can recognize the main components of the processes that give rise to the complex world of intertwined aims. In this task, I will use the standard language of economics, in which the utility function maps circumstances to the individual’s preferred outcomes. Nonetheless, I will situate the discussion in its appropriate context,

recognizing that certain environments will only support certain kinds of individuals, and that we can in this way conceive of endogenous utility functions.

We could thus think of the utility function as the relatively unrestricted happiness of the agent, in absence of upper level restrictions. In a prisoner's dilemma, for example, the utility function is represented by the proposed payoffs presented to the prisoners, while the behavioral function is represented by what the individual actually *does*. Moreover, the constraints to which the behavioral function has to abide by are different than those presented by a typical budget constraint, as they literally define the agent, in a sense compelling the individual to 'like' consumption bundles that he would otherwise object. One could only dread to wonder, for example, what a suicide bomber feels as he presses the deadly button, but it must be nonetheless something that he would not have felt in the absence of social pressures. The proposed behavioral function describes at all times how the agent makes decisions in a setting that may most often consider the effects of her actions on future selves, offspring, family, other members of her species, and all sorts of lower and upper level layers of which she is composed and to which she belongs; in contrast to a more abstract utility function that pertains to her individual, most immediate self.

Aside from a utilitarian description of people, the hierarchical agent representation must always reflect the relativity of the definition of upper and lower levels, as they can only be classified as such in reference to an agent; they represent a conceptualization for classifying and disentangling complex decision-making and behavior. However, if we think of processes in terms of the concept of work—where the concept refers to the constrained use of energy—then the lower level is better conceptualized as the free energy or fuel, and the upper level represents the constraints. In this sense, the upper level may not necessarily be thought of as a restriction of possibilities, but rather as a market in which the isolated agent can trade behavior for some form of retribution, perhaps corresponding agreeable behavior of a peer.

Thus, the superstructure presents a set of options for what the individual can be, making some placements of the utility function 'cheaper' than others. In this sense, individuals could become unfriendly, criminals, traitors, and generally undesirable, but the society will create mechanisms to discourage such outcomes, here modeled by an 'expensive' region in the realm of action of the agent's utility function. By the same token, honorability, courage, honesty, friendliness, and all sorts of likeable characteristics will be praised and encouraged, thus modeled as 'cheap' regions. Moreover, agents may encounter temptations they avoid because of previous experiences in which the then current self became uncontrollable and detrimental to the welfare of the more general self.

The realm of action of the utility function can be constantly modified by time, society (i.e. any given upper level), and peers. A society may relatively consciously decide to close its ranks, follow rules, establish norms, or stand united so that the resulting, more efficient group can more effectively compete with peer groups. For example, one can separate American and Mexican societies, as representatives of the developed and developing world, according to the narrowness of the realm of action they present to their citizens. Whereas in the U.S. littering, tax avoiding, unpatriotic behavior, and general rule-breaking are frowned upon or even prosecuted, in Mexico, people brag about their untamable nature, about their independence from the system and their corresponding lack of cooperation with upper levels. This differentiated and path-dependant pair of outcomes is most likely a function of the awareness individuals have of the higher self (in this case, the nation); American society has invested substantial resources in educating citizens about their place in the world.

The socialization structures that normally shape individuals significantly limit the realm of possibilities for the utility functions that are the product of a person's upbringing. Parents go out of their way to ensure ample resources and norms, so that their children grow up to become adults that please them; while later on in life, the individual is commonly constrained by a family of her own. In a society stressing this sort of family networking, individuals disenfranchised from such opportunities may be faced with the harsh reality of relative isolation.

If we think of realms of action as open markets, we can conceive of them as boundless and simply having regions of near zero probability, (such as a lazy bee,) or unaffordable choices (such as becoming a serial killer.) When we want to understand the influence upper levels have on an individual's decision process,

we must consider the realms of action that each level above her presents, which can be thought of as optimizable functions. If we solve these optimization problems and present the solution to the individual decision maker, a mapping function could be constructed to reflect all the influences of the upper levels. The individual may perceive this influence as a unitary force, conceptualizing it as ‘the force of collectivity’, and considering it in a tradeoff with pure utility. As pictured in Figure 5, the individual chooses a point along this possibilities frontier; a choice that is represented by a maximization of the individual’s behavioral function.

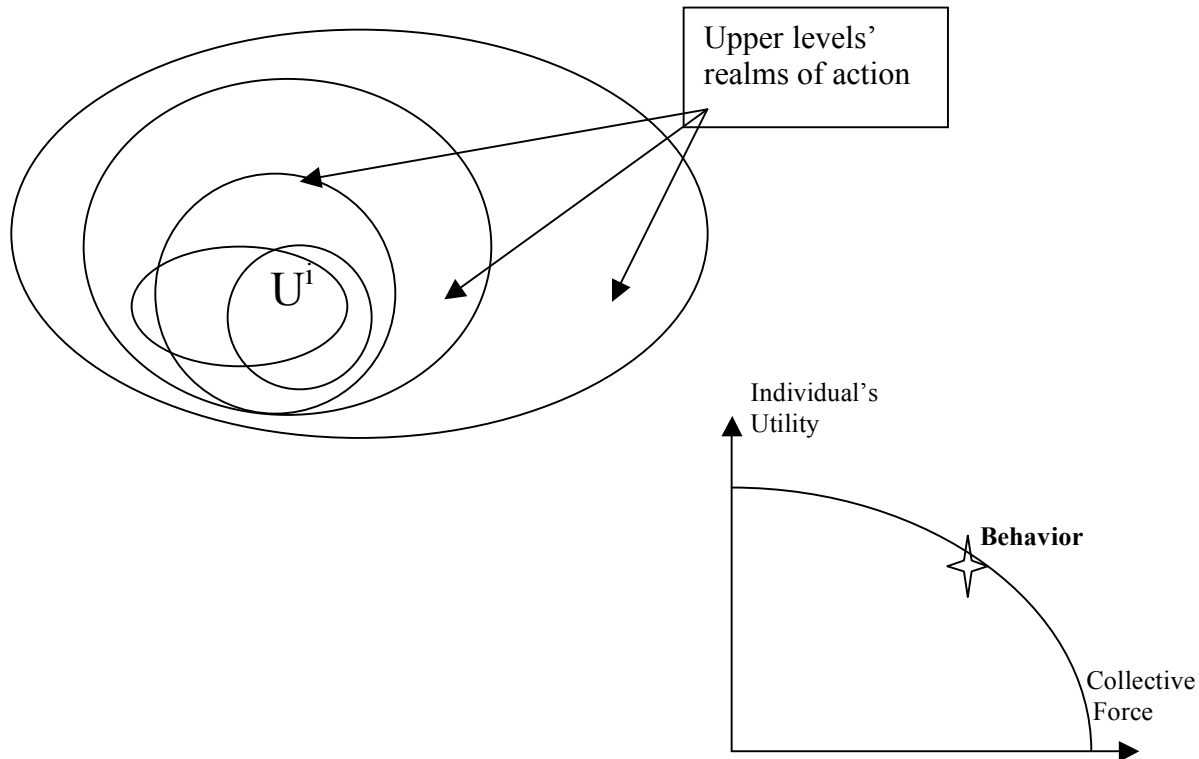


Figure 5

An agent will be constrained by upper levels or implicit rules of behavior that will not let its current self make decisions that are in complete disagreement with the set of organizations to which it belongs. The upper levels may encompass one inside the other, like a babushka doll, but may also be defined according to different criteria. The actual behavior of the agent will thus recognize the tradeoff between the objectives of the most current and narrowly defined self and the objectives of the more encompassing upper levels.

The agent may not recognize that the upper levels are multiple, but instead may conceive of them as an environment in which it defines itself. The objective function of the agent represents its force of “individuality,” and can be modeled in an emergent utility function that would include the influences of lower-level composition. Once again, the agent represents a level that is commonly composed of a series of subagents engaged in negotiation, cooperation, and competition. Their resulting force, however, has a correspondence in its utility representation, which may be transitory. Such an agent does not act in isolation, as the upper levels it belongs to impose their muscle in its decisions, through forces that can be translated into the same dimensions as the ones in which the utility is defined. Any given hierarchical agent thus chooses between these two attractors in varying degrees, and according to their individual strength. The perfect agent to describe the perspective I am proposing is a group of colluded oligopolists, such as the oil-producing countries of OPEC. As a whole, they share the common interest of jointly behaving like a monopoly and restricting their production, but they cannot avoid having an incentive to deviate and produce above their quota. I believe the proposed framework of hierarchical agents can be broadly applied in agent-based simulations for describing complex interactions among players defined in different dimensions. In the final two sections of this paper, I discuss contexts that accept the applications of the

proposed model, as well as some other insights that could be gained by looking at the world through the proposed perspective.

3.- US vs. THEM.

Suppose we are interested in understanding the reason why Latin American countries are relatively poor, and in the process of creating an agent-based simulation to analyze the problem, we find ourselves lost as to how to define our agents. Should we have a macro view in which we make the Latin American bloc an agent that deals with its position in world politics? Or should we consider each one of the participating countries as separate entities? Should we rather adopt a micro model that starts from individual people? One of the main ideas of this work is to bring attention to the way in which a more general model would deal with the connection between the different dimensions of analysis. In the history of economic thought in the twentieth century, for example, a push was made for developing macroeconomic models with strong microeconomic foundations, but the need for tractability left the discipline with no choice but to resort to linear models of interaction, as well as unrealistic assumptions about rationality. The era of agent-based modeling, however, can forgo linearity and perfect foresight, and take a leap towards a more general and realistic model that can describe the emergent strategic interactions at each relevant level.

Suppose, again, that in the construction process we begin with Mexico as the objective agent we are trying to model. As figure 1 shows, such an agent does not have clearly defined borders. First of all, not all aspects of Mexico can be defined as lying inside the concept of Latin America. In particular, the Mexican business cycle is now much more related to Canada and the U.S., its partners in the North American Free Trade Agreement (NAFTA), who among many other issues may have a strong incentive not to let their southern neighbor fall into the recurring Latin American financial crashes. Moreover, Mexican lawmakers are currently close to reaching an agreement in which Mexican nationals living in the United States would be able to vote in the upcoming Mexican presidential elections. Under this agreement, it should not be long until this group has representatives in both chambers of Mexican congress. Official figures reflect that the ranks of Mexicans Americans and Mexican nationals living in the United States has reached more than 18 million, or about a sixth of the total Mexican population.

Aside from the issues entailed in exploring the many layers of agency involved here, the Mexican example sheds light on the fact that no agent we will describe under the proposed paradigm will have clear borders delineating what belongs to it from what does not. Consequently, any agent considered will be “fuzzy” in the sense that it will be described according to many different criteria, that it will have relatively blurry borders, and that it may or may not have overlap with other agents depending on the criteria we consider for defining it. Correspondingly, when we think of agents at a micro level, we think of individuals who are to some degree the one they are in the very moment we look at them, but also to some degree a reflection of a longer time horizon of the self; to some degree the relationships they are involved in; as well as to some degree a piece of their family, of their community, and of their nation. I believe that such a view would not be too controversial, but requires that we reconsider what we mean by the pervasive concepts of selfishness and altruism, since they are not clearly defined when we consider such fuzzy agents. The view I propose builds upon W. D. Hamilton’s Kin Selection theory, in which altruistic acts for closely related individuals may be considered selfish behavior. In addition to the situation of shared genes, I would like to extend this concept to explore to other motivations for cooperation or confrontation. This question in turn leads to the need to differentiate to which extent another human represents a potential mate or potential for the advantages of mutualism, and to what extent it represents potential competition.

The most important point to make here is that defining where the self ends is not an easy task. In particular, and as Marvin Minsky’s research points out, people may not be the unbreakable wholes philosophers conceived of in the past, but are rather a collection of many selves entrenched in intense negotiations. If we could isolate and talk to one of these many internal subselves, (and I believe we do, most of the time) at some point, what we would hear about is a connection of many different relatives with very strong interconnections, and about some upper-level description in which most of them form part of a ‘we.’ An internal agent who likes ice cream may settle for going to the movies—the train of thought goes on to

say—not because it likes the movies, but because the agent that does is a relative, and therefore to some extent another version of the self.

Another example could ask people how much they would be willing to pay for a sibling to obtain a thousand dollars, or perhaps to give it to an anonymous person living in the same city, or to an anonymous person living in another country. If the person asked is willing to pay something, then that act can be considered altruistic, but I propose giving consideration to the possibility that humans consider other humans themselves to some degree, and in this sense paying one penny for someone to receive a thousand dollars is a bargain for our selfish spirits. Vernon Smith won the 2002 Nobel Prize in economics for “having established laboratory experiments as a tool in empirical economic analysis, especially in the study of alternative market mechanisms.” One of his most important experiments was that in which he had people playing anonymous prisoner’s dilemma games. Given the setup and payoff structure of the game, the rational dominant strategy is not to cooperate. The actual participants, however, generally opted for cooperation. As I will discuss in further detail, I propose the idea that humans generally think of other humans as partly themselves.

I propose a research agenda aimed at elaborating a model that could describe the world according to the degree to which observed behavior reflects agents that do not completely act as individuals, but rather act or think of themselves as interrelated, and therefore working for similar objectives. The idea is that we could look at the way organisms interact in a given population and establish a measure, or coefficient of relatedness. This concept would be similar to the proposition of kin selection theory, in which organisms can be thought of as wholes, but also as one half their siblings, one eighth their cousins, and so on, according to their degree of relatedness. The goal would be to read and describe behavior in order to define corresponding coefficients of implied relatedness among populations, societies, and cultures according to similar principles. At this point, the proposed model would not have any hypothesis to test, but would rather serve as a classifying mechanism. The model would thus be able to classify left and right wing politics depending on how large they define their ‘we’, with the former generally having a more encompassing definition than the latter.

Aside from an unclear definition of what is to be considered ‘I’ or ‘we’, we must of course see the other side of the coin: at some point we leave the ‘we’ territory to enter into the world of ‘them’. Returning to the example of asking people how much they would be willing to pay for a person to receive a thousand dollars, let us now imagine adjusting the question to specify the recipient as a citizen of a Muslim country. Given current political circumstances, my belief is that many in the United States would be willing to pay for this game not to occur, because they would fear that some of that money could end up in the hands of a terrorist: the enemy. The far-reaching aim of the proposed model would thus be to better understand the way in which the boundaries between the ‘us’ and ‘them’ are drawn, for they are immensely important, relatively variable, and often misunderstood.

3.1 BORDERED MAXIMIZATION.

In this section, I present several issues I find critical in understanding the way ‘us’ vs. ‘them’ boundaries are established, and they can be roughly categorized under the banners of *cohesion* and *bordered maximization*. Bordered maximization refers to the process by which the size of an agent is determined, while cohesion refers to the optimal design of the hierarchical agent. At the aggregate level, cohesion does not imply that all lower-level actors are acting in coordination with respect to each other, but rather that an upper level exists for whom the actions, design, and general topology of the lower-level agents maximize an objective function. What this implies is that an upper level may want some of its lower-level agents to compete, and in this sense perfect competition could be ideal for the upper level that an economy represents. Complete cohesion thus exists when—given a set of constraints, perhaps of the nature of the subcomponents—there are no reorganization possibilities that will better serve an upper level, and it is therefore a relative concept. As for bordered maximization, the proposed insight is that a subset of the components of an agglomerate will attempt to maximize what we can define as their objective, or utility function, in a process that can be considered selfish or myopic from the point of view of the whole, but optimal for the acting subset. Moreover, and joining these two concepts, the ruling subset involved in a

bordered maximization process will likely have its internal cohesion as one of the most vital characteristics for achieving organizational success.

The framework outlined above thus creates a free-for-all in the proposed world of multi-dimensional hierarchical agents. Again, aside from running simulations in which results are obtained for a particular setup, the proposed model would serve as a classifying mechanism for recognizing the strength that emergent agents have in a proposed area of study. As I will discuss in section 4, there may be no general definitions of optimality in a hierarchical world, only ones that are relative to an agent in a given level. Nonetheless, I will propose a benchmark position in which all observed behavior is optimal, so long as we understand the active agents involved in the observed behavior, always taking into account that the relevant agent may be abstractly defined.

Under the proposed frame of work, albeit contained within the paradigm of selfish behavior, I rethink the definition of an individual, so that we can re-conceptualize such things as selfishness, altruism, and rationality in a representation that recognizes the inherent subsequent decomposition of all agents. Literature on the matter has focused on broadly defining altruism as actions that negatively affect the individual who performs them while at the same time benefiting someone else. This notwithstanding, it is commonly shown in the proposed models that altruistic cooperators can survive in the end, mainly because their resulting stronger groups can defeat other groups formed by non-cooperators. With that result in mind, we can think of less myopic agents, or some created by eons of evolution, which will be forced to give up their lower-level objectives in return for upper-level ones. The true altruism of nature is to act today to guarantee the existence of a tomorrow.

Bordered Maximization is thus the term I would like to coin for describing the process through which an agent defines the line that separates itself from the rest of the world. It is a dual process, as it decides the size of the agent at the same time that it decides its identity. Imagine a view of the world in which we begin with autonomous atomic agents that form increasingly larger and more complicated structures. The structures they form are fragile, and find themselves in a constant process of reorganization. What is it that a newly formed agent maximizes? Its objective function. But why not forgo a myopic definition of the self, and create a coalition? Why not break apart from the current coalition and form a group whose identity lies closer to the core identity of the main actors in the resulting groups? In a most general setting, the answer to these questions is naturally beyond the scope of a single paper, but I nonetheless classify the general type of processes involved.

Bordered maximization represents the process or force that drives the reorganization of hierarchical agents. It is a reaction to opportunities for niche-filling and defense against ever deadlier enemies. In this last sense, a growing enemy may provide incentives for two otherwise independent parties to join and form a coalition. One could then think of America's unparalleled power as triggering the formation of a stronger European Union. A similar rationale would posit that Germany delayed its natural incentive to unify until the twentieth century because of a lack of a common enemy. A potential enemy also promotes cohesion, just as we saw when the American people stood united after the attacks of September 11, 2001.

Nonetheless, even as this structure formation takes place, smaller regions may have an incentive to break apart. The European Union has American power as an incentive to unite, and is willing to make an investment in incorporating poorer Eastern European nations into their conglomerate, but the richest countries in the region, such as Norway and Switzerland, can gather enough individualistic strength to warrant their exclusion in the union. The bordered maximization of Switzerland's objective function merits its exclusion from the upper-level political entity, in part because the benefits of the union have not risen enough to reach the 'inclusion reservation price' of these wealthiest of nations in the region. A related issue refers to the definition of right and left political platforms, where the position of the party reflects the intention of including a specific portion of the population. Models forecasting the placement of these lines would be extremely valuable, and the proposed framework can serve as a language in which we can express their meaning.

The electoral process in a simple majority-based democracy is a prime example of a bordered maximization of objectives, in which the winning party does its best to find an objective that satisfies the desires of a

social contraction attempting to win the election, and in order to do so, they must bring the majority of the population to their side. The process is naturally a very complex one, and an experienced politician may declare that no situation is ever the same, but we nonetheless may explore common strategies that the optimization should include. In choosing issues and sides, a political party may stand on principle, express opinions based on an ideology, or take on the issues that affect a previously defined group of people. In the process of expanding its base, the party may also decide to propose wedge issues that will allow others to join the coalition and decide certain aspects of public policy. As such, the party's platform must take into account the reaction the electorate will have to each one of their proposals.

Bordered maximization has two main implications for implementation: on the one hand, the political party searches for positions that are most agreeable for its core identity, while on the other, it must be careful to choose the issues that bring to its side the necessary amount of people to win the election. Naturally, these two sides are interwoven, as finding an acceptable political position for the base must include a directive not to alienate the general population by doing so. Many inside America's two parties would love for their representatives to be much more out of the mainstream, but that would imply becoming extremists in the generalized political scale. Moreover, in looking for wedge issues that are external to the base (in the sense that the party's base does not feel too strongly about them, or that they are willing to concede them,) the platform must look for issues that will bring the largest amount of people, while protecting the core issues of its base. In a heterogeneous population full of potentially controversial issues, I would expect a two-party political arena of intense competition to produce very tight elections, and multi-party ones to provide platforms that will serve the needs of all large definable groups. Groups of people who are concentrated on a subset of the main issues—and who will correspondingly join the party that best addresses those issues without asking for much on other issues—will be highly coveted and powerful.

In a simpler two-party system, how will bordered maximization find successful coalitions? How will elasticities of becoming a member of a party be analyzed to form a winning ticket? To provide an idea of the answers, suppose we were to enumerate all possible combinations of all conceivable wedge issues. Without any regard for which issues tend to clump together for ideological or cultural reasons, random division of the issues must by definition create winning majorities, albeit very odd ones. With the passage of time, however, I would expect things to fall into more understandable combinations in which the winning platform has some sort of internal consistency and is potentially more cohesive. Inside the winning coalition platform, I would expect to find a group that is cohesive enough to accept the full implications of a proposed issue agglomeration, and that is formed by people who find themselves better off here than with the competition, or who belong to a group that is benefited inside the platform. The issues themselves chosen will tend to be ones that provide large subjective benefits to the group (or provide large electoral gains) without alienating others too much. In addition, the winning party must take into account that new elections will come again later, and that its actions must not disturb the losing party too much, for the disenfranchised could become a liability to the ruling coalition, as the opposition could let go of most of its issues in order to concentrate on just winning, and in extreme cases could start a revolution. This effect brings to light the fact that the entity created by a bordered maximization will at all times find its position in a longer-term upper level of existence, with an objective function that does not get trapped in immediacy.

The previous example is a simplification of the larger description of the determinants of group formation, of the interplay of forces that gives rise to the delineation of the 'us-them' boundaries. In the election example, the short-term goal is well established: to win. The more general concept of bordered maximization can be applied to a wider set of issues where the objectives are not as clear. Such would be the case if the political party not only wants to win, but wants to create an enduring majority platform. Establishing the general framework in which models to address these questions can be created must harmonize advancements in economics, sociology, evolutionary biology, political science, psychology and social psychology within the greater paradigm of complex adaptive systems.

3.2 IDENTIFICATION: LIVING IN A WORLD OF EMERGENCE.

Who are the agents in a hierarchically decomposed world? How do they think of themselves? Instead of considering the actual thought processes happening inside humans and other animals, we shall focus on what their behavior implies about whose objective function they are maximizing. As I have argued, organisms are not independent from the evolutionary progression of their species, both historically and looking into the future. The structures of long-term optimality appear to the organism as instincts, ‘natural’ drives, needs, feelings, etc. To a human, these upper levels may have even more abstract undertones, and come about as rules of thumb, superstitions, morals, conscience, implicit rules of social behavior, and the like. Among other things, these tunnels of behavior are epiphenomena of the social structures that represent the scale of exploitation of the agent. Sexual animals are incomplete. Not that they immediately die in isolation, but as such would not form part of the recursive process of its species. Social animals are thus smaller pieces of a puzzle. The nature of atomic decision makers is one intrinsically possessing an appetite for abstract resources that superstructures can provide. Therefore, the individual will normally face a payoff structure that allows him to survive if he stays with the proposed boundaries of the higher level.

Aside from other incentives for joint effort (such as specialization and risk sharing), individuals mostly make decisions based on incomplete information, and often without a clear objective to maximize. Therefore, the current individual may often find herself in a position where she may evade searching for drastic, egoistic outcomes, and instead settle for strategies that have proven useful at maintaining the long-run individual alive, healthy, wealthy and the like. In the same fashion, groups that have found cooperative structures to abide by will tend to survive longer than groups that are penalizing themselves internally for lack of such coordination.

I thus expect the payoff structure of the standard prisoner’s dilemma to be relatively hard to conceive of for a realistic human, since most current selves and individual agents will naturally be afraid of—or compelled by evolution to avoid—strategies that overemphasize rampant immediacy. The behavioral function exemplified by this reaction is perfectly compatible with the survival of a higher version of the self. The true benefits of cooperation are themselves the emergent property of an upper level in the hierarchy, and one which we cannot understand if we zoom in too close to the individual’s decision without understanding her within the context of the superstructure to which she belongs. Moreover, the setup of a one-time deal given to an individual with no recourse to an uncertain future is unrealistic, as organisms who think like this would shortly perish without intertemporal cooperation.

Any activity one could think of, from walking to talking to tool-making, requires some degree of coordination, and for activities with longer-term horizons intertemporal coordination is indispensable. Like humans, animals living in an environment with seasons must also coordinate thrifty summer selves with needy winter partners. The upper level to which I am referring works as a process of identification with intertemporal selves as well as with other related individuals in a group. Disconnected organisms are therefore a near oxymoron. I claim that there must be an evolutionary force connecting otherwise autarkic selves and individuals, and that the process of creating an ‘über-self’ or a society could possibly be modeled as the creation of a market in which intertemporal participants and peers can negotiate on coordinated strategies. Such a process may be imperceptible from the localized perspectives of the individuals involved; however, from an evolutionary point of view, the reorganization process may happen at all levels and with far-reaching implications.

Examples of human identification with higher-level social structures are pervasive, and there is a significant body of literature studying these issues from cultural, political and sociological perspectives. We can find examples of it in our everyday life and in the news. One such shocking instance was the case of rampant suicide among Brazilian soccer fans when the Brazilian National Soccer team lost the final game of the world cup in 1950. Other examples include the solidarity members of the same religion feel towards each other, or the tendency of waiters to leave higher tips when dining out, as they relate to peers in these relatively anonymous settings.

Another incentive for individuals in a group to merge their incentives lies in the fact that discontented individuals have the tendency to be nasty. The payoffs of a more realistic version of the prisoner’s dilemma

are naturally determined by the history of the game—most importantly, whether the participants are anonymous or known participants, and the previous experiences of the participants. Non-cooperative actions can always set a precedent for something that will later affect an agent directly related to the defector, as cooperation and lack of it are normally persistent. An angry participant may thus set drastic reactionary payoffs, possibly to instill fear in possible rivals of future similarly afflicted selves, just as the grim reaper strategist would drastically react to engagement. Imagine a world in which agents face repeated prisoner's dilemmas against anonymous and perhaps different partner every time. Furthermore, this world is inhabited by modified Tit-for-Tat agents, whose strategy is to cooperate or not according to what the previous partner did. The projected effect of a deceived Tit-for-Tat player in this network of similar agents implies that a decision not to cooperate will set off a chain reaction with negative effects that a finite group can always return—possibly heated up through the complexities of what we could refer to as the game's karma. After all, any malicious action played against a peer can become institutionalized and potentially harmful for the meta-temporal self. Naturally, just as with reciprocal altruism, the degree to which these effects will be relevant depends on the size of the group we define as 'us.'

4.- OPTIMALITY IN A HIERARCHICALLY DECOMPOSED WORLD.

4.1 A HIERARCHICAL MODEL OF EVOLUTION.

The discussion in this paper has attempted to establish that the self does not have a clear boundary. This idea has immediate repercussions for the way we think about natural selection and the evolution of organisms. Although the approach has been developed along the lines of economic science, I propose that a model developed from these ideas could just as well apply to evolutionary biology. I believe that the boundaries of current research in both disciplines are facing similar obstacles, and that they share a similar historical paradigm. At their core, both of these disciplines are based on a selfish and unitary agent. Both paradigms claim that all aggregate complexity can be traced back to the lower level of the system: the strategies and actions of the selfish agent. Along this line of research, individual selection theorists such as Richard Dawkins conceive of a world full of evolutionary stable strategies. These population-wide Nash equilibria at the population level refer to situations in which observed patterns in aggregate behavior are simply strategic best responses of the individuals to an environment full of potential competitors. I have made similar arguments in this work about the linear models prevailing in economics. In summary, and in the language developed here, both research agendas de-emphasize the existence of any level other than that of the individual. As I understand it, the idea of emergence reflects the fact that different and irreducible levels of interaction will naturally arise in complex systems such as the ones studied by these disciplines.

The agent, as I have defined it in this work, is both subject and object. It is a combination of levels of interaction. Organisms are thus a product both of the struggles of their ancestors, but also the reflection of an environment that begs to be exploited. They are therefore part exogenous and part endogenous to Mother Nature. It is through this lens that I would like to consider humans, who will partly be independent creatures, possessors of free will, but who are also partly created by an array of upper levels that 'suggest' agreeable utility functions. I believe this conception stems directly from the concept of complexity, in which wholes are more than the sum of their parts. If we believe in that proposition, then we should expect to find a world full of emergent phenomena, with distinctive levels of interaction that have agency of their own. If these levels exist, it follows that they must evolve.

Evolution is therefore not a process in which autarkic individuals compete to be the fittest, but rather a historical multiplex of intertwined levels. In this sense, the phrase "survival of the fittest" takes on a more specific meaning; it does not only mean that the genes of the best-suited individuals in a population will survive, but rather that the populations themselves will become more adaptable, deadlier, more complex, and all together 'more fit' at exploiting environments. To clarify this point and to make a personal contribution to the timeworn tradition of using sports metaphors to defend an idea on evolution, let us conceive of natural selection as a process which seeks to assemble a Super Bowl-conquering football team. What is important to recognize is that the rules of the game are themselves being selected, and it is in this sense that we can talk about more fit individuals, but also about more fit selection processes. The main

purpose of this discussion shall reflect an attempt at defining different levels of optimality, and the claim is that we can elaborate scales in which some groups (perhaps as large as a whole species) can be deemed better fit for long-term survival than others or than other versions of themselves, in a similar way in which we can compare individuals and categorize some of their traits as more fit than others.

What is special about the organisms that inhabit today's earth? For one, they reproduce in one way or another. There is the possibility that in the past, a highly advanced species decided not to do so, but we would have to hope for paleontologists to let us know of their curiously ephemeral existence. This is one restriction that applies to any existing species. Countless other restrictions apply to each specific species, without which their populations would not survive long. Since these restrictions are external to the individual organism, they may be modeled as upper levels. Natural Selection applies, then, not only to individuals, but just as well to groups, laws of behavior, species, and groups of species. Because there is no visionary designer looking out for the correct implementation and discovery of these upper levels, it is plausible that a group as a whole will behave in ways that do not maximize its potential to keep its recursive process going; however, those that do happen to follow strategies optimized with a longer term horizon are more likely to be observed. This idea bears some resemblance to the popular Murphy's Law, whereby, for example, we can say that you will most likely find the slowest lane in traffic, if only because when you do, you have a lot of time to think about it.

To provide one example of an upper level that makes a species more fit for long-term survival, we could consider a species of ferocious animals in which the males fight for courting rights. On opposite sides of the world there are two versions of this species, with corresponding populations that are identical in all respects except one. The aggressive version has fights that go on until one of the parties is dead, while the other version stops a fight after a few acrobatic movements that clearly reveal which one of the two contenders would prevail in a deadly fight. Although the odds for surviving the battle in the lethal species favor the stronger individual, it is often the case that *both* fighters leave the field mortally wounded. In the second, more peaceful version of the species, the defeated male allows the dominant one the space he asks for. Unless the fights serve as a necessary means for population control, I claim that the peaceful version will have better odds at long-term survival, and can therefore be declared fit. The fact that a daredevil type could infect the population and spread its genes should not be a sufficient theoretical basis for precluding the appearance of these cooperating groups. The individuals of this less aggressive species will have upper levels that preclude them from fighting to the death.

Organisms thus represent the current stage of long, recursive processes that have searched vast phase spaces to find the combination of complicit levels that sustain the information they entail. Random mutations are generally not an option that will normally be part of such processes; rather, controlled ways of adapting to an ever changing environment will naturally develop. In other words, exchanging a known successful situation for an alternative that will be drawn from an infinite space of unknown possibilities is unfeasible. For example, the emergence of sexuality provides a systematic way for the mixing of populations with information that has proven useful and constantly recreates the organism in order to stay ahead of predators and parasites. We could think of an enhanced sexuality that provides other improvements to the population, such as superior genetic material that is phenotypic and recognizable. The benefits of sexuality can come from different levels; at a localized and instinctive level, in which mates with compatible genetic structures might produce smells that create attraction, or at the less-localized level, with non-genetic social norms that lead to correct population control.

The focus of this essay is humans, but the general idea of this discussion should be applicable to any species. The important point here is that humans are alive because they are fit to be so in innumerable ways. In line with the theme of this work, what I mean is that a human's utility function is not created in a vacuum, but is instead the product of a culture and millions of years of natural evolution. One of the many reasons why humans have been so successful is the fact that we learned how to cooperate with each other. In this perception, the space of possibilities before the advent of Humans included many different possibilities for an emerging species to exploit. If we optimize that space in search of what kind of species would be best fit for such an environment, what we would find is that a cooperative one would be better equipped than its non-cooperative counterpart. In this sense, the optimized phase space of the environment represents a series of demands for exploitation, while the existing species represent a series of supplies.

When the two meet, a recursive and self-sustaining entity is produced. In this sense, we can think of evolutionary pressure as a situation in which one group of individuals is in position to mutate and fill up a niche that the environment presents. Social structures must then be analyzed in the same fashion we would think about the limb of an organism: the way in which evolution acts upon a social network must be studied just as an ethologist studies animal behavior. In other words, evolutionary pressures shall act upon individuals and subcomponents of an individual, but just as well upon the structures these individuals form.

4.2 DEFINING AN OPTIMIZING AGENT.

Vernon Smith won the Nobel Prize in Economics for bringing the science closer to an experimental environment. In one of his most famous experiments, he had a group of students play repeated and anonymous prisoner's dilemmas. As it has now become standard, the game presents both players with a payoff matrix with one dominant strategy: not to cooperate. If both players play rationally, their best response is to defect on each other. However, the result of Smith's experiment was a surprising amount of cooperation. This outcome is particularly interesting when we consider that the empirical results of most studies in economic science fall below the expectations of the theory, as in the case of the recurrent appearance of irrational speculative bubbles in the stock markets of the world. In this case, the theory says that—if rational—Smith's class should have made a given amount of money: instead, they made much more. The implication is that a lack of rationality leads to a better result. A host of explanations have been provided in the literature, but what does the conception presented here have to say about this? A general description of such situations must include the corresponding consideration of a superstructure at all levels; the individual or group making a decision may bow to an upper level that commands him, whether he is aware of it or not. Taking an all-out present decision is commonly harmful, understandably fear-provoking, and eventually in contest with a long term survival plan.

I have argued that the payoff structure that a player faces in the prisoner's dilemma refers to its abstract utility function, which can only be understood in isolation. The payoff must not be the sole source of information for the decision maker. I believe that the participants in experimental games such as the one discussed above are making decisions as representatives of calculating potential intertemporal selves, or of the human race, and not as isolated individuals. One of the most important reasons that humans have been so successful on this earth is because they learned to take advantage of mutual cooperation, of the possibility of social synergies for finding better scales of environmental exploitation. As such, a human makes decisions under the constraints of its incomplete nature. The agent at play is a long-standing evolutionary force that has allowed humans to conquer the world. In order to develop a model that thoroughly explains cooperative behavior, we must head in the right direction, and the claim is that we must look for these blurry upper levels that are explicitly defining behavior; binding the individualistic nature of the prisoner to a strategy that does not make sense when analyzed in isolation. Naturally, each level influencing behavior may be defined in very different dimensions. This can be seen in the time frames over which these levels change; a person will often modify her behavior immediately in response to a stimulus, while upper-level social structures generally change over long periods of time.

In terms of rationality as it is most commonly defined, one can think of the actions of benevolent cooperators as irrational, since they do not follow their best interest. However, as these cooperators generally end up prevailing (if only thanks to the conditions of the model), one may wonder if an agent in more realistic circumstances may be aware of such outcomes—as well as of the likely retaliations of deceived players—and be therefore acting in a perfectly rational way, once she acknowledges the better information of peers' characteristics, retaliation functions and gains of cohesion, and in that sense develops a less myopic mind. In comparing intrapersonal choice and personal decisions, it is interesting to note that an individual who acts in a resolute manner (that is, forcing through a precommitment a future self to act in a manner which is then not optimal for her) is deemed by some as irrational; a statement which is in direct contrast with calling myopic someone who does not recognize the opportunities of intrapersonal rearrangements to achieve preferred outcomes, when having a longer horizon to base their decisions.

We can think about Vernon Smith's proposed total payoff as an environment. Cooperative as they were, the humans that participated in this experiment were able to better exploit their environment by extracting more

money from the researcher than they would have if they did not cooperate with each other. By acting as they did, they were able to explore the possibilities of hyper-productivity provided by an upper level. How can the upper level, this cooperating group, protect itself from infection by the non-cooperating type? That remains an open question. But Mother Nature is savvy, and through the ages has figured out ways in which these upper levels can become more cohesive, can be binding to their subcomponents. Socially, primates have long found ways in which upper levels can be enforced, as there is conclusive evidence showing that several of our close species enforce social norms, even to the detriment of the individual. Humans have found even more creative ways to build strong superorganism levels, and as such we can understand the advent of global religions. Christianity, for example, in part became so successful and was adopted by the Roman empire because it allowed for the creation of a large human network, in what can be considered the first ever wave of globalization. From the point of view of the role it serves as an upper level, this religion shares some important characteristics with the modern religions that have the most adherents: it forgoes localized, idiosyncratic gods, and instead places humans as equals in the face of divinity; it calls for cooperation, among at least a subset of (believing) humans; it provides a set of rules of behavior that facilitate the integration of larger social groups, creating a sort of ‘credit bureau’ in which followers of the religion can find peers and trade with them increased and valuable information. Islam, on the other hand, also defines the concept of *jihad*, or a coordinated attack on an outward group.

In light of the coalescing power of this collective force, should the expansion of these coordinating structures continue indefinitely? Hardly. The force of individuality will not take long to strike back. The tradeoff between these two forces reflects the outcome of a bordered maximization, and the particularities of each situation will determine where the boundaries of the levels appear. Nonetheless, as the history of the world’s game unfolds, evolutionary pressures will solicit larger, more complex and adaptive organisms or networks of coordinated organisms that find themselves better at exploiting the ever changing environments offered. These larger organisms will have a nature, a design that has optimized the possibility for exploitation *at that level*. The ‘agglomerate individual’ will have to become cohesive, and organize its subcomponents in order to maximize its emergent utility function, since suboptimal internal coalitions can materialize, in the way a drug addict cannot overcome the desires of a strong drug-craving self, or a tyrant may force a society to extinction.

It is this tradeoff that researchers such as Samuel Bowles capture when they talk about “The Co-Evolution of Love and Hate”, as well as other related models where individuals have an incentive not to cooperate, while the group as a whole has an incentive to have cooperators that will allow it to prevail against other, less effective groups. I have made the analogy of such a tradeoff to the one faced by colluded oligopolists, which have a group incentive to behave like a monopoly, but cannot avoid the temptation of breaking with the agreement to sell a larger amount of overpriced product. I find this concept to be the kernel of a new social theory of complexity that attempts to get at the root of the interplay between individuality and collectivity. This dichotomy is one of the most important issues in the process of building a new complexity paradigm for the social and biological sciences. In terms of evolution, for example, one may wonder what kinds of species a world would produce? For example, would it produce macro-organisms whose subcomponents have very little leeway for adaptation, or a myriad of microorganisms that can easily adapt, but have very little group coordination?

As Bowles models it, the groups encounter incentives for internal cooperation as they prepare for battle with other groups facing similar circumstances. The battle against these peer groups can be seen as a representative example of the implications of a Darwinian evolutionary environment. More generally, we do not need such structured competition to realize that internally cooperating groups will tend to survive longer and possibly grow, as their wealth, health, and general well being will be higher than those afflicted with internal conflict. Just as the process of evolution perfects individuals, so does it have the same effect in groups and societies. The surviving members of a cooperating group, however, will not be ‘fittest’ at an individual level; their individual traits and natures make sense only within the context of the cooperating group. In this sense, while social structures facing competition may be forced to recombine in order to guarantee their continued existence, the internal biology of a living organism—with its structure determined by DNA—is usually impossible to reorganize within a generation. Following Bowles’ definition of the ‘assumption of efficient design’, we can think of living organisms as smoothly coordinated societies, as their subcomponents ‘magically’ aid each other, quasi-perfectly specialize in tasks, and have

well-developed methods of internal communication (such as a central nervous system). Also, as opposed to known problems of principal-agent relationships, organisms are most commonly indivisible, as all the subcomponents go ‘on the same boat’, since the whole must worry about the slightest cut, which could potentially mean death for the whole organism.

If we extrapolate the previous discussion to the organization of an army, whether it is a modern one or one of antiquity, a human one or one formed by ants, it should be clear that evolutionary forces will shape increasingly deadly bellicose groups in a process that may be sped up by the increasing returns of victory and the hegemony of the victorious party. Similarly, the construction of a winning political platform is a process that joins overlapping issues until it creates an identity that can appeal to more than fifty percent of the electorate. Although it may reach across many different coordinates of concerns, the identity, platform, ideology or political party will most commonly have the incentive to present itself as internally consistent and relatively homogeneous, as its internal cohesion will likely prove key to its success. One striking difference between the evolutionary processes shaping the internal organization of a living organism and those in charge of reorganizing companies, armies and political parties is that—unless we believe in some form of external and perhaps divinely unfolding design—the former is the product of a natural elimination process, while the latter is the product of an accelerated, conscious human attempt at ‘staying ahead of the circumstances’. I will explore this issue in further depth in the next section.

How do we know where the agent’s boundaries will appear? In other words, in light of the discussed complexities, is a tractable model possible? Perhaps the real-world, complex and chaotic nature of the elements at play will prove too extreme for predictions to be made. Nonetheless, we can begin by observing where these boundaries appear in reality, and then use backwards induction to portray the forces at play that could have given rise to the observed behaviors and structures. The proposition of this paper is that we can begin with a benchmark position in which all behavior is optimal, so long as we identify the actual agent that is enforcing its will, and that such will or behavior may actually be the resulting force of the interactions of several abstractly defined agents. In this sense, when we arbitrarily zoom in and analyze a relatively well-defined agent, we may classify its behavior as suboptimal in relationship to its own abstract objective function, but only because we would be artificially studying it in isolation; the upper level may force lower-level members into behaviors that are only optimal for the former, and it is in this sense that we can understand the behavior of a soldier. Traditionally, we have begun with a clearly defined agent and tried to understand its actions as a maximization of objectives given constraints. In the proposed paradigm, we assume maximization occurs, and then work towards the delineation of the benefited entity involved.

Just as Plato argued that no man ever knowingly does wrong, an agent who does not choose its maximizing strategy in a prisoner’s dilemma is obviously choosing based on something it deems better, even if she has been programmed to do so, whether by a precommitment or an evolutionary influence. In other words, agents may be the product of an evolutionary process that rewarded cooperators, such that agents of current generations are naturally born ‘good’. In this sense, the better-than rational results of Vernon Smith’s experiments may be understood by appropriately characterizing the acting agent: in this case, the people who decided to cooperate against their incentive to defect were not acting as individuals, but as representations of the higher level of evolved and cooperating persons. In this way, the upper level asserted its agency. Individuals are born into structures that are usually in many ways set for them, with incentives for cooperation and trade that are inescapable for the individual, as their nature intrinsically possesses an appetite for the resources that only the superstructure can provide. Other examples of this process include the behavior of working ants, unmonitored employees committed to the success of a company that is not their own, or deadly suicide bombers acting in the name of their religion.

Furthermore, I believe that a parsimonious representation could synthesize all the influences upper and lower levels have on the individual—in the same fashion that a utility function synthesizes the aims of all its conflicting subcomponents—with two forces describing any level, or decision situation. The first one refers to the most immediate ‘agent-self’, with the composed objectives of its lower-level subdivisions, while the second one represents the collectivity the agent belongs to. While the former is selfish and myopic, the latter forms structures, both in space and in time, possibly exploiting ‘reorganization margins’ in a long-term evolutionary process.

With the perception of humans as conglomerations of decision-making units in mind, I propose analyzing groups as structures that can be thought of as possessing predictable behavioral patterns, analogous to the proposed emergent behavioral function of a person. Any group which builds upon networks of lower-level collections of agents may be thought of and modeled as having ‘a mind of its own’, as well as relative control over the subcomponents constrained in their corresponding realms of action. In this sense, the agent can be thought of as a person, a family, a social class, a political party, a country at war, or a species as a whole. One of the main goals of this proposition would be to create a scale in which to measure a group’s welfare, arguably just as the standard utility function does for a set of disjointed individuals. I explore this proposition in the last section of this manuscript.

4.3 BETTER THAN PARETO.

The discussion so far has left us with an important absolute principle: all possible definitions of optimality in a hierarchically decomposed world are relative. The relativity is a direct result of the hierarchical nature of the system, where each agent binds the subagents that compose it, and is bounded by the super-agent to which it belongs. Therefore, defining optimality implies defining a level, for it is only by placing it within its appropriate context that we can describe how one behavior is better than another. For example, an evolutionary process of competition—where each competitor could potentially be better off without an enemy—may be bound by the optimality of encompassing, more efficient groups or social superstructures with a larger scale of exploitation. The unfolding of history then implores the reduction of the system’s entropy, as the continuous reorganization of actors aspiring for immortality leaves no room for improvement other than in unexplored realms of complexity.

Theoretically, Pareto optimality has been used in neoclassical economics as the wall that separates positive from normative issues. Basically, a group or society finds itself in a Pareto optimum when there are no possible reorganizations in which everyone is either benefited or left indifferent with the change. Based on this principle, the main welfare theorem of economics states that, under certain conditions, a trading collection of agents leaves no ‘leftovers to spare’—or, in other words, arrives at a Pareto optimum—when all agents can pursue individual objectives. Correspondingly, this view declares redistributing income to be an intrinsically normative issue, since no matter how beneficial it may be for some, it would necessarily have a negative effect on someone else in the process. Thus the rupture with Pareto, since there are no markets in which benefited individuals can compensate affected ones. Implicitly, the concept of Pareto optimality adopts the individual as the level in which it operates; any improvement that can be made may not pierce through the boundary that defines a person. As I have argued, however, there is no clear definition of what should constitute a person.

A revised analog of Pareto optimality may be developed. The concept would have to recognize its relative nature, as it is therefore defined only in relation to a particular level. In this sense, the new concept of optimality would represent the internal aspect of bordered maximization, or an all-encompassing reorganization effort to find the arrangement that best suits the agent, or level in question. In relation with standard Pareto optimality, the new concept must recognize all the emergent abstract agents of a super level. In this way, Pareto optimality precludes income reorganization, whereas the proposed concept would look at all possible markets, coordination and interrelationships among the agents, so that, for example, a reorganization could be accepted if every citizen of the country has a decent chance of obtaining a higher income in the future, or if everyone will have someone close them benefited by the new arrangement.

Even then, there may be some who will not want to proceed with the rearrangement. Ruling parties, ruling classes, and empires will have much to lose in most reorganization initiatives, and will therefore be against them, even if the changes could improve long-term general welfare or survival probabilities of a group that includes them. This is once again why the proposed enhanced version of optimality must be considered relative; we may never find a concept, a level that encapsulates every member of a group. Nonetheless, we can just as well declare our level to be a target representing, for example, the poorest 99% of a nation’s population. Because we have grown used to democratic ideals, most of us could accept a new arrangement that improves the lives of the majority of a country’s citizens, since they themselves can always elect a political platform that represents their interests. What to do with the minority? Once again, the ruling

majority must recognize that the decisions it is making are to be placed under an umbrella, long-term upper level, and should therefore not break the rules of behavior that have brought them here. In other (more drastic) words, the ruling majority should definitively not kill off the minority, as this action would likely institute new rules of violence that would eventually negatively affect those who could seemingly be benefited today. In summary, the optimization process of a group must not be seen in isolation, since it is itself inside an upper level that represents the interconnections that define the group within a broader, longer-term perspective.

Relationships described by upper levels can therefore be thought of as abstract markets. In this sense, a parent may gladly give their lives to save their children, and a rich person may willingly pay high taxes to sponsor programs that reduce street violence. Pareto optimality has been useful to describe situations in which we cannot conceive of a market where affected parties of a social reorganization can negotiate on a settlement. For this reason, rupture with the Pareto optimum has been used as the argument that makes active income redistribution a normative issue, since there are no realistic markets in which benefited individuals can compensate negatively affected ones. Looking inside a partitioned individual, we find a host of potential markets in which feelings, regrets, a sense of accomplishment, or flat-out pain serve as methods of compensation and retaliation; in which, for example, an affected subset will have to be heard and recompensed in order to avoid conflict. In contrast, standard economic science implicitly states that a person is either a unit (even when seen in an intertemporal sense), or that it is subdivided but with perfect internal markets, where any change in the external situation is completely balanced among the internal subsets. In both cases, Pareto's guidelines have no dealings inside an individual. Once we accept that one cannot break apart the threads that make up an individual, however, the atomic version of society implicit in Pareto optimality—upon which the main economic paradigm is built—is rendered artificial.

Most importantly, given natural appearances of in-group cohesion and out-group hostility, affiliation to 'Pareto laws' could doom a subset of agents who abide by it, forcing them into a myopic view that cannot envision the possible collective gains of cooperation in the face of inter-group conflict, if only because the reorganization may hurt some of its members. Furthermore, the agents in a group must find an optimal design for working together, stressing the benefits of heterogeneity, specialization, risk sharing and the like. As the research of Herbert Simon points out, many economic structures do not organize themselves in a market scheme, but rather optimize their underlying internal relationships with intricate structures of command. Thus, the company is in a constant reorganization process to create a structure that maximizes profits (its objective function or an aggregate measure of 'fitness'), and then turns around to allocate the spoils among its participating members. In forming the superstructure, participating subsets will find themselves as part of the upper level in a process of identification. The disciplines of political and social psychology have devoted significant research to understanding how humans form social structures. In particular, postmodern discourse theory proposes the idea of subject positions, where individuals play a number of changing roles brought about by the ideologies to which they subscribe.

As we relate identification and optimal (read: deadlier) scales of social groups, one of the most important aspects that define the formation of groups is organized violence. Agents are naturally worried about hurting peers without backing from a *self-sustaining* group, because the future of the game is naturally unpredictable. This refers back to the concept of *optimal scale* for the agent, in the sense that the necessary group size is dependent on peer circumstances and auto-sufficiency constraints. Specifically, the social structure involved in conflict must procure the creation of 'altruistic' and risk-loving individuals, and be large enough to allow itself to lose some heroic sons. These expendable individuals would seem to forgo maximizing their utility functions, but only if we fail to recognize the 'cage' they are being put in by the superorganism to which they belong. In other words, meta-temporal structures in conflict will often need to be able to support subsets that take on strategies that are not sub-game perfect. It is in this sense that we can think of the development of religion as an upper level that may become so strong that the subagents completely let go of their individuality, as is the case of the dreaded suicide bomber. I explore some reasons why we would expect the appearance of violence in the next subsection.

4.4 THE BUSINESS OF VIOLENCE.

I have argued extensively that a social level can be modeled as an agent, in the sense that it can be ascribed an objective function and—given an appropriate time frame—be expected to maximize it by reorganizing its subcomponents. This framework was proposed, for example, to describe the advent of cooperating groups that better exploit their environment, which may involve other, less cooperative groups. The question remains as to where we place violence in this framework. Before embarking on such a dark discussion, I must disclaim that I am attempting to describe an objective view in which customary moral perceptions of violence are dropped. Just as Howard Bloom suggests in “The Lucifer Principle”, this is a view that recognizes violence as pervasive, and therefore seeks to find its unfortunate roots. In line with the proposed principle of conceiving all behavior as optimal, so long as we correctly define the agent whose function is being maximized, we can conceive of the behavior of the suicide bomber as optimal for the group of people in conflict who support him. Nonetheless, I would like to follow up on that horrible assertion by hopefully asking if we can expect the appearance of levels even higher up in the hierarchy for which this violence is suboptimal. In other more specific words, can we expect the appearance of a United Nations superorganism that will enforce the elimination of such practices, perhaps by even more violent means?

In this view, I maintain the perception of relative optimality created by these overlapping hierarchical agents, and search for the roots of the agent for whom the violent act is optimal. Throughout this essay, I have given examples of such behavior, and it may prove convenient to remind us of some of them to analyze under the proposed conception. I have talked about an animal species for which deadly fights that determine rights to female access may be considered suboptimal, in the sense that both parties could benefit from avoiding the fight, if only enough information correctly assessing who would win it could be exchanged in advance. A brother could fight to death to avenge an attacked sister, in an act that may be considered suboptimal unless we look high enough in the ladder to find an agent—such as the family’s honor—for which the strategy is optimal. Such an upper-level strategy may have provided defense to the group in many occasions, but requires occasional individual action to remain a credible defensive threat.

These examples are chosen to express the relativity of violence’s beneficial and detrimental aspects. A world with no violence could be considered best, but given the realities of our nature, a second-best world may have to include police forces and righteous wars. To define if violence is optimal for a given group, what we need to ask is if there is an upper level that can present the hyper-productivity benefits of peace and allocate it to all parties involved so as to avoid conflict. In other words, as a group prepares for war, it may originally believe that it is in its best interest to do so, but it may nonetheless one day recognize that, once all the secondary effects of war are accounted for, a path of peace could be preferable. The path of peace would only make full sense to them if initiating the aggression will lead to enough misery for the group and related future generations to outweigh the potential spoils of victory. It is in this sense that I expect Mother Nature to create ‘hawks’, predators, species devoted to war, in the same sense that we would always expect the appearance of non-cooperators within the cooperating group, of free riders that will take advantage of the opportunity for exploitation of the peaceful chickens.

Violence may be optimal for the upper level representing a group if the resulting winning parties are correspondingly better at exploiting the environment, or if the environment is insufficient to support all individual members of the group in conflict. In this sense, the example of the species where the males fight to death could be reconsidered, as the fights themselves could be Nature’s way of finding and perfecting the right kind of environment exploiter, and the deadliness of the fights may be reflecting an appropriate means for population control. Along a similar vein, it is a long-standing tradition in economic theory to praise the benefits of competition, as this social structure allows for efficient outcomes that maximize both the quantity as well as the quality of goods produced by the economy.

The last point brings us back to the important concept of optimal scales of exploitation. As I have argued, all organisms are incomplete, for they only represent a link in the long recursive process that gave them life, and that forces them to maintain it. To have a meaningful life in the evolutionary sense, sexual animals are forced to find a partner and mate. Many species including all mammals are even further constrained by the fact that the younglings need constant care as they grow up. In a social species, where the optimal scale

of environment exploitation requires the coordination of more than one individual, a peer may possess aspects of both a competitor and a potential partner. The optimality of the upper level correspondingly implies subdividing the agents until it finds cohesive subcomponents that are cooperating or competing in nature. This phenomenon appears in treaties such as the Geneva Convention, where countries that expected to have further bloody wars accepted common laws that would benefit the global upper level.

I conclude this discussion by stating that further research studying the sources of conflict is warranted. In the current war against terror, for example, my belief is that unless the western world decides to engage in a full-out war that would redefine the concept of a holocaust, the sources of terrorism will indefinitely continue to produce fighters, for the social enemy is an upper level which is fed by precisely the violence that intends to suppress it, creating a never-ending recursive loop. On the other hand, unless you come from a place where corruption is unbridled or where the state is solely a force of repression, most of us think of police—the ultimate representative of violence’s force—as beneficial. War in an atomic world with defined political boundaries, however, has many of the representative traces of a myopic maximization, even in the face of the currently purported ‘expansion of freedom’ theme.

The final section of this manuscript develops the optimality theme of the present section, but attempts to expand it to other realms of social organization, exploring the effects of social domination within a society, as well as the implicit pecking order established among nations in the unfolding process of globalization. As a capstone for the proposed hierarchical methodology, I encapsulate the discussion by conceiving of groups and social structures as agents that maximize objective functions by reorganizing their individual institutions, in the same way that a person orders its internal subagents in an emergent utility function. In other words, the ultimate goal of this model would be to create a scale in which to measure a society’s welfare, arguably just as the utility function does for a set of disjointed individuals.

4.5 BEYOND PARETO: THE BASIS FOR A SOCIAL WELFARE FUNCTION.

Pareto Optimality reflects a conception of the world where society has reached a plateau, or as economists would say: one in which all margins have been exploited. This directly stems from the neoclassical economics’ idea of how an efficient market emerges in a free, capitalist environment. Given the initial endowments, such a market extracts all possible benefits society has to offer, and any attempt to improve society beyond this point must include a reallocation of resources among individuals, which is considered a value judgment. As I have argued, I believe that general improvements beyond Pareto are possible in the realities of social evolution, but neoclassical economics has decided to deem them subjective. We must find ways in which we could measure the benefits of social structures such as those devised in the modern constitutions of a democratic nation state. The optimizing mechanism of design need not necessarily be one ruled by a conscious higher level, but may instead be constructed by the complexity of the localized decisions of every member, who in every decision propose strategies that may be genetically imprinted, selectively neutral, or brought to a wedge by evolution, to a margin that serves a social purpose.

Arguments for relative social equality and elimination of extreme forms of dominance are considered subjective by neoclassical economics. On the other hand, I consider them to be relative but nonetheless objective. The discussion should turn to pinpointing what the relevant actors are, and I think there is very strong evidence supporting the idea that humans are not to be thought of as individuals, but rather as highly interconnected social creatures. The individual should therefore not be the basis for the paradigms that describe humankind. Communism, on the other hand, seems just as arbitrary to me, for there is also insurmountable evidence of Man’s greed and selfishness. The nature of the environment for which groups of humans compete does not only involve procreating, but also creating niches of privilege that beg for competition. Any social structure we could conceive of, and certainly any that would arise naturally, would be an upper level which benefits from competition, for only then can it continue to breed talented people and organizations to fill its necessary components. Classical economics does have an indisputable truth when praising competition. For that matter, imagining a living world without natural selection would challenge even the most creative of science-fiction authors (but not so your average creationist).

The point at hand is that discussing different social structures in terms of optimality is an objective exercise, so long as we are clear about what the optimizing target agent is, as well as the definition of the currency we are using for measuring the improvements. The search for candidates to be that target should lead us to long-term, upper-level, more natural agents that will keep the recursive process of life that brought us here intact, if not producing more fit individuals. The fact that the present displays the largest cooperative human masses ever assembled encourages me to believe that the optimal scale of environment exploitation includes most of the world's population; in other words, that the force of collectivity that brings nations together makes larger cooperative structures more fit, and therefore represents an algorithm that may eventually create a global government body that will prevent large scale wars, as well as a regulating body that will internalize the externalities of nature's destruction.

Wilfredo Pareto himself conceived of a society in transition that would naturally reorganize until it reached that equilibrium at which further reorganization efforts would have to hurt someone to benefit anyone else. There is thus no reason why Pareto-improving reorganizations should not be made, since by definition there would be no individuals harmed by such adjustments. However, if the environment includes groups of individuals, societies, and nations, then these upper levels must also be taken into account in issues of optimality. The motive for focusing solely on individuals when making these assertions is to avoid intrapersonal comparison, for only utility functions are defined objectively in the classical paradigm. In contrast with that position, I argue that groups can be thought of as having objective functions, just as disjointed individuals do.

Positive and normative economics have formerly been clearly divided by the wall that Pareto optimality represents, but the concept of a Pareto optimum could doom a subset of agents who abide by it, forcing them into a myopic view that cannot envision the possible collective gains of cooperation in the face of inter-group conflict. Therefore, the Pareto optimum can no longer be used as the threshold that divides normative from positive economics, as it may be perfectly rational for a group to reorganize itself, maybe hurting some of its members in the short run, but maximizing long-term welfare. Moreover, if we accept the proposed notion of individuals that are not as clearly defined, then the concept of Pareto optimality becomes unclear. For example, if a social rearrangement makes everyone better in the future, but hurts someone's present, can we say that it was Pareto improving? How about if it improves the lives of all our children (granted all members of society have children)? Even more perversely, if wealth is relative, then someone in a society may want everyone to be poorer, does Pareto optimality then imply that we should keep everyone poor so as to not disturb such an aristocrat?

Unless we refine the concept, Pareto optimality is not a good measure of general welfare, but simply a protective barrier for the status quo. It represents the way in which the ideas of capitalism, laissez-faire, and general economic liberalism have been theoretically cemented. This notwithstanding, only strict capitalist zealots would call for the end of the progressive tax system. Full liberalization of economies leads to monopolies and cartels. It also leads to an underproduction of beneficial public goods, as well as private goods with positive externalities such as education and health care, while overproducing goods with negative externalities such as pollution. In order to counteract these negative effects, the upper level that society represents will take control, internalize the externalities, and take the necessary actions for general improvement to occur. This occurs because a majority of the agents recognizes the generalized benefits in their own individual realities, and trades in markets in which alternative behaviors are exchanged, thereby bringing social classes together into a coordinated social unit. For example, a rich factory owner may realize the benefit of decreasing crime rates and improved workforces that better income distribution and improved education may generally bring about. It is precisely for this reason that anti-globalization zealots may have a good point when they claim that globalization often amounts to exploitation. Globalization differs from modern nationalized capitalism (in which negative effects have been naturally internalized and partially corrected,) because it acts by interconnecting distant regions with no recourse to each other. The result is that many people are not impacted by the negative externalities of production processes that take place thousands of miles from their national borders. In other words, our nature may possess a tendency for the exploitation of humans by humans, driving society towards slavery if we sleep at the steering wheel.

A common point made in international trade theory is that trade is good for the economy. Opening a previously autarkic economy represents a vast reorganization of a society in which—at least in the short

run—some are benefited while others are hurt. The claim that trade is good is deemed an objective one because it can be shown that the monetary gains of those benefited is sufficient to compensate those who are negatively affected. The fact that this claim can be proven theoretically is no evidence that this is actually the case in practice. In the real world, producers of those goods that are no longer competitive in the open economy are seriously harmed. Before free trade agreements are signed, these people clearly recognize this fact, and commonly protest accordingly. Nevertheless, the theoretical argument can be made that trade is good because we can define a monetary market in which affected agents can trade gains and losses (if we imagine that some people who lose their jobs will be harmed in psychological ways that cannot be monetarily quantified, then we would be back at lack of interpersonal comparability.)

In order to develop a group objective function, we must then conceive of other kinds of markets in which agents of all levels can exchange norms and behavior. Internally, the superorganism presents a realm of action for the utility function that improves the individual's position from isolation, but likely promotes behavior in accordance with the group's benefit. Moreover, if the formation of a social structure requires rupture with Pareto optimality, the potential benefits of that structure themselves represent a force that will encourage its appearance, forcing the reorganization to coalesce, even if it is unfair to the status quo. Every demand creates its own supply.

We may thus consider a social welfare function that allows the citizens of a country to accept social changes and circumvent Pareto optimality, if only the citizens have enough connections to a potentially better future. Of all possible reorganization schemes, we conclude by focusing on the highly important issue of income distribution. Although it does represent a Pareto optimum, I will assume most of us would agree that a society where one (selfish) individual has all available income is not optimal, and will likely be the victim of a hostile takeover or an internal revolution. The fact is that a society in such peril will likely change through the reactions of its individuals, who will mainly act based on local information and reflecting their personal interests. As individuals involved in such revolution may find themselves with an 'interdependence perspective', in which they see their success as highly positively correlated with the success of their close peers, we may expect the appearance of cohesive groups. Particularly in the last half-century, sociologists have made significant progress in explaining different human perspectives on the origins of in-group cohesion and out-group hostility, but I have not found in these theories the definition of group welfare I search for. Could humans bring forth a set of general rules that lay the basis for a social welfare function? I claim that the civilized world has been implicitly doing so for millennia, through religion, accepted social norms, state nations, democratic ideas, and culture.

I will conclude now with a couple of examples of such processes: The American experiment and the Mexican suboptimal social structure. The American constitution represents one of the greatest experiments in the history of conscious human evolution, and the scope of its results are written every day in the history of the most powerful social object the world has ever seen. It now seems obvious that the Founding Fathers had a great practical vision for how to organize a society with such great potential, as it had a fresh start with a relatively equal society (of its free men) as well as vast resources, both natural and human, as they inherited the rich cultural and scientific tradition of European nations. With the grand ideas of the French revolution, backed by the condensed knowledge of the Western world, the Founding Fathers applied all their scientific and pragmatic wisdom into a constitution that separates powers very effectively, with an executive branch that remains relatively external to the government, and a commander in chief with significant control over bellicose decisions; a legislative power designed to 'listen' more intently to the needs of its heterogeneous constituency, and is somewhat comparable to a living organism's nervous system; and a judicial power that ensures the permanency of rules that should remain relatively constant, providing fertile ground for projects that need longer time horizons to develop fully, and makes sure that rules are applied equally to all individuals, thereby maintaining the hard fought equality of the citizens that will eventually make the country strongest, with reliable institutions, networks, generalized habits and the corresponding social norms.

As the analysis of the proposed model suggests, one of the most important aspects that the separation of powers is that it allows for the social organism to better discern the many different aspects of change. For example, and continuing the analogy with the proposed psychology of a disjointed person, we would expect a social object to be 'fearful' of change, as it is hard to digest all the implications of an organizational

scheme that has never been tested, and that affects citizens through so many different and regularly unquantifiable dimensions. The separation of powers represents a mechanism by which individuals can discern and disentangle the all-encompassing effects of social change, knowing that their government will follow rules of engagement that will attempt to protect their need for security and the exploitation of known benign structures and norms, while accepting the possibility of change and further progress through the exploration of other, perhaps socio-culturally neutral, dimensions of the social conundrum. Aside from the original design, the fact that the American system has been bipartisan has made for a government that is most ready for immediate action, a trait most important in the art of inter-group conflict (i.e. war), since the ruling coalition of a multi-partisan political system could partition, particularly in times of political stress.

In terms of what the American political system does to reach a social organization equilibrium, we could see the electoral college election system as one which eliminates information when the federal system does not 'get to hear' whether the state likes a particular candidate overwhelmingly over the other, or if instead the race was really close. The process creates the so called "battleground states" where the candidates need to focus, and are therefore left to localize the federal decisions. We could also imagine how the composition of a person's internal decision system could produce a comparable objective function, with a central decision unit (or a receptor in its neural network) that only hears the direction of the feeling that the subsection of selves presents, instead of its intensity. This could be a natural mitigation mechanism against extreme desires. For example, some one who falls into a drug addiction may also think at all times that he is doing what is in his current best interest. He is, however, not recognizing the all-encompassing effects of his actions upon related selves that have been subdued by the tyranny of his drug craver selves. This is irrational just as committing suicide would be, especially for someone who gets up in the middle of the night ready to assassinate all other sleeping selves. By analogy, the American electorate elections system could frustrate rogue, out of the mainstream states from taking control.

In contrast to this conscious and successful plan, social systems such as the one Mexico has present a political ruling class for which it is convenient to have a significant part of the electorate left in the dark, and this can be seen in the vague—and mostly devoid of controversial issues—electoral campaigns that their candidates run. The most important example is that no significant attack on generalized poverty is proposed, and this is the direct reflection of the fact that the system is implicitly colluded with the economic ruling classes, in an agreement that precludes an evolution that could bring about the birth of a more equitable state. Would such a state be preferable? I believe we could easily measure the improvement in any reasonable dimension, and definitively provide visible monetary results after a generation. A system that provided real divisive issues in the ballot could thus be considered more desirable, but we must then keep in mind that the society should keep track of not creating a society that changes too fast; a society that leans too much on the side of *exploration*, and then forgets about the *exploitation* it could do of proven to work arrangements. The attainment of appropriate middle points is, once again, brought about by such things as separation of powers and constitutions that have a long-run vision.

Vision is something that the Mexican nation has unfortunately lacked throughout its history. Some may be surprised to hear, for example, that in the eighteen century Mexico and other Latin American countries had per capita GDPs that were comparable to that of the United States, and I believe that a model based on the proposed paradigm of this paper could account for the most important reasons why Latin America's economies have sunk since, while the U.S. has flourished. Many people in South America believe that the reason why their countries are poor is because they have been enslaved by the western world into a situation in which they only produce raw materials that have seen their relative prices constantly decline. This notion is most commonly known as dependency theory. I claim, however, that the problems of Latin America are internal; their societies, as a working upper level, are simply ill-conceived. The regulations, institutions, implicit rules of behavior, and other norms that play such important roles in coordinating developed societies are extremely weak in these countries, whereas citizens of rich countries are proud to stay within the sensible boundaries that their societies propose. Latin Americans commonly brag about every rule they break, and boast their seemingly untamable nature. In countries such as Mexico the rule of law hangs by a delicate thread, forcing social norms to take on the tasks that a proper state should operate. Furthermore, the ruling class of the country seems to lack the desire and certainly the urgency for improving the impoverished condition that defines most citizens. In contrast to the fact that wealthy individuals in the U.S. have to work hard to stay in the upper class, the Mexican rich are born into an

aristocratic system that allows them to be rich for generations, even when they are incompetent at managing their wealth and generating economic expansion.

Dependency theory calls for a disengagement of the Latin American countries from a world that has bound them to poverty. It is hard to imagine how this strategy would prove useful in an increasingly globalized economy. I agree with this position in one respect: the historical reasons why the Latin American block is poor do stem from the tragedy that colonization represented for it. The imperialist social structure brought about by this period created a polarized and racist society that inhibits development. I disagree, however, in thinking that the solution would involve less involvement from the outside world. On the contrary, these countries have a lot to learn from the developed world. They need to understand how to model their societies with a broader perspective in mind. The colonial lock-in societies of these countries in fact restrain and decelerate any process of economic development, since a few fortunate individuals may thwart a process of general improvement that negatively affects them.

What can a country like Mexico do to leave this trap? I optimistically believe that advancements in our understanding of complex adaptive systems shall bring relative prosperity when correctly applied by policy makers to solve the multifaceted problems that developing countries face. In the Mexican case, any solution will have to start from the bottom up, as this may be the only way in which the society will adopt long-lasting solutions. In particular, the necessary tax reform that the country needs must be made at a localized level, because citizens that have been previously deceived by corrupt governments will not accept a federalized taxation system whose benefits are hard to perceive; instead, they would much prefer to see the direct improvements that paying their local taxes bring. The creation of a national social network must then start at a grassroots level, where the interconnections between the involved agents are still strong.

In summary, we can think of social objects as having a behavioral function in analogy to the more understandable reality of a disjointed individual. In this sense, the idea is that the complexity of a social object will have predictable ways in which it will respond to any given situation, just as a person's behavioral function does. To this end, we can compare an electorate to the proposed myriad of selves inside the individual, which can form coalitions and have debates for the upper level to take the decision that suits them best. This could be modeled as a neural network in which receptor neurons receive the person's partitions wishes. The receptors can be in turn anthropomorphized as a system of extremely capable political analysts who calculate all the effects of any particular decision, determining whether or not it would be approved by the particular political system in place.

The proposed welfare function will take into account the nature and corresponding complexities of all the levels of agents it encompasses, quantifying such abstract notions as social cohesiveness, identification, and play allowed for the individual's utility functions. The proposed welfare function would then provide a scale in which we can evaluate alternative social structures. Most directly, the proposed function could map social structures to corresponding future changes in per capita GDP. One interesting research avenue to seek this elusive concept may involve understanding the economic system in terms of the physics concept of work: the constrained application of energy. If we think of a society as producing work (perhaps measured by per capita GDP), then we must understand the right constraints that will maximize output. This optimization process will likely bring us to a tradeoff having too lax of a system in one extreme (as in the case of the Mexican society), to one that is too restrictive in the other (as is the case in a communist planned economy). Both of these systems do not promote the creation of niches to exploit, and neither do they encourage the appearance of the driven workers and entrepreneurs that constitute the combustible of the economic machine. The objective function must then find the correct structure design that brings the society to the edge of chaos in which production, innovation, the environment, trade, and peace can flourish.