

ECONOMICS WITH A BIOLOGICAL FOUNDATION

ALDO RUSTICHINI

ABSTRACT. We argue in three lectures that economics should develop on a foundation of biology. We first define what economics with biological foundation is, its position in the history of the discipline, and why this development is necessary for economics to remain a meaningful science.

We then outline the three main directions of development for such new science of economics (neuroeconomics, genetics and personality theory), and illustrate research in these three fields.

(Aldo Rustichini) DEPARTMENT OF ECONOMICS, UNIVERSITY OF MINNESOTA, 1925
4TH STREET SOUTH 4-101, HANSON HALL, HANSON HALL, MINNEAPOLIS, MN, 55455
Email address: aldo.rustichini@gmail.com

Date: July 17, 2023.

Acknowledgments This paper is based on a group of lectures (with the same title as the paper) I gave at the Winter School of 2022, organized by the *Centre for Development Economics*, the *Delhi School of Economics* and *The Econometric Society*. I thank professor Soumendu Sarkar for inviting me and for extremely helpful comments given on the first draft. I am also thankful to the audience in the lectures for patiently listening, inquiring, and probing.

Keywords: Neuroeconomics, Genetics and Economics, Personality Theory.

JEL: D9.

Competing Interest: The author declares no competing interest.

Funding: No funding was received for conducting this study.

1. INTRODUCTION

I present here my general view of what economics with a biological foundation should be, and outline some of the transformations of the discipline that are necessary if one adopts this point of view. At the conclusion of this essay, a reader may wonder whether the discipline described here is economics as we usually understand it. Indeed, it is not economics as it is today; but I think this is how economics should be in the future.

If I may be even more bold, if economics doesn't look like this in twenty years, it will not be useful anymore. Economics as a science has changed considerably since the late 70's. Most of the changes have been for the worse. Now we are at a crossroad, and we need to evaluate how to proceed. Very likely, there is not going to be a unanimous consent on the way forward, because in large part the transformation in the sciences (particularly the social sciences) is in part the outcome of deeper changes in how the power in society is managed: economics is by necessity part of an wider ideological struggle.

Thus, the future developments in economics are not just theoretical issues, that we need to discuss as scientists that are interested in furthering a neutral discipline. Most important, policy issues that we discuss daily depend on the answers to the questions that we analyze today.

This essay has three chapters. In the first one, I will give a broad overview of the topic. In the second and the third I will go more in detail on two of the topics that I briefly described in the introduction.

But let us first define our object of inquiry.

1.1. What is Economics with a biological foundation? In its simplest terms, economics with a biological foundation is the science that proposes to understand economic and strategic behavior, mostly in humans, by taking as the foundation of this understanding the biological determinants of this behavior. It is therefore, as first step, an operation of boundary delimitation; it begins with the statement that if we do not include a set of biological facts into the analysis, then our efforts to explain and predict economic facts will not be just useless, but would create and further a false understanding, which is worse than no understanding at all.

“Include into the analysis” has to be taken here in a strong sense. It does not mean that economists should borrow a few scattered tools developed in other sciences, use them in the occasional task of interpreting the phenomenon at hand, and then turn to other tools when the next phenomenon appears. It will not be enough to go to the next building on campus, pick a few ideas in an “interdisciplinary effort” and go back to the economics building with a few psychology papers in the memory of our computers. “Including” means that the biological determinants of economic and strategic behavior should be part of the building itself, laid as we lay the other bricks of the structure.

We have tried so far to give a definition. But to best understand what economics with a biological foundation is (or, more appropriately, should be, because the process of building it is just beginning), we should locate it in the historical process of development of economic theory.

After the end of the classical school, economics tried, with the marginalist school first and with the neoclassical synthesis later, to establish a methodological foundation on the basis of simple psychological regularities, such as the decreasing marginal utility. The program since its outset was based on the hope that a stripped down psychological theory would suffice to provide the foundation for economic analysis. According to this program, looking for conditions defining an equilibrium, or just first order conditions for an optimal solution, was all that was needed.

Some familiarity with the writings of the founding fathers shows that there was always awareness that this initial setup was a convenient simplification, and that perhaps later developments should make the theory richer. The opportunity of this enrichment never came, because the critique that this foundation was insufficient was instead presented (in the late 70's) as the radical criticism that this foundation was fundamentally wrong. Economics with a biological foundation is the attempt to fill this gap.

1.2. Two views. The conceptual basis of neoclassical economics is the simplifying but insightful assumption that people are fundamentally rational and selfish, and that the process described by the mathematical formulation of the corresponding maximization problem is also, in first approximation, a descriptively good model of the process that they follow in choosing. Economics has until recently proceeded reducing the problem of behavior to an optimization problem, and using the mathematical solution of the problem as a prediction on behavior. As we mentioned, economics has also acknowledged that this assumption is a simplification; however, it has not yet ventured into the task of replacing this simplification with a realistic model. By refusing to do so, it has exposed its flank to the easy criticism of ignoring deviations from the rational behavior. The observation that the theory is an approximation has been transformed into the damning criticism that the theory is useless because it is false.

To be more specific, neo-classical economics has two fundamental insights. If we abandon any of the two, economics will be useless. The first is the *individualistic idea*, that the social outcomes are best explained as the outcome of the interaction among individuals who try to achieve what is best for each. The second insight is that each individual is *substantially selfish*, and also reasonably good at deciding what is good for him (or her), so the interests he pursues are very clearly his own.

This individual behavior is then combined into a social behavior through two equilibrium concepts, that is, the general equilibrium in market economies and the Nash equilibrium for strategic situations. These two cornerstones

of economics as we know it are the most ferociously criticized, particularly in recent years, for example in behavioral economics.

In our analysis below we add a third component, what I call “biological materials”, which includes the genetic basis and the neural basis of individual behavior.

1.3. An opposing view. An opposing view to the one we just described assumes instead that human nature does not exist as an independent entity, separately from society, but it is simply the outcome of social processes. This general thesis has three corollaries:

(a) Human nature is infinitely malleable, and thus subject to the beneficial transforming action of the reason. The role of an intellectual (in particular of an economist) is to point to the right direction, so reason may eventually prevail. Reshaping human nature is the substance of progress. And the march of progress will be faster if the reason can take advantage of the enforcing power of a state.

(b) Men are all potentially equal in all features that are relevant for social life: they have the same skills, in all dimensions, in the same quantity.

(c) Every systematic difference between groups that appears in statistical analysis (for example, between the mean and variances of some characteristics) is either false or due to social factors. These factors may be, for example, privilege for some or discrimination for others. Only social factors produce these differences.

Such statements are the political consequence of the theoretical assumption that human nature in itself does not exist. There maybe other ways to justify them, but denying the existence of human nature is the speediest.

This is a variant of idealism, and you can find it in unexpected places. In some of his most famous passages, Marx proposes the following idea: all the evil in society is coming from the division of labor which is induced by capitalistic mode of production. This thesis is a constant theme through his entire life: from the preliminary drafts of his youth (Marx and Engels (1845)), to the later works (Marx (1875)). He proposes the thesis he proposes that the division of labor has created a society where each one of us has a narrow sphere of activity, as a hunter, fisherman, herdsman, or a philosophical critic. This classification does not occur because, by his own nature, a man is more inclined to be a hunter or a herdsman or the critic of philosophical critic, but because society puts each in an arbitrarily assigned role. Marx then envisions (Marx and Engels (1845), “Private Property and Communism”) a future society where it will be

possible for me to do one thing today and another tomorrow,
to hunt in the morning, fish in the afternoon, rear cattle in
the evening, criticize after dinner, just as I have a mind.

This hardly seems what he really aspired to do. If you know a little the biography of Karl Marx, you will also know that he never went hunting, never went fishing, never reared cattle, but constantly criticized. If you could

put Karl Marx in the new world after the full establishment of communism, he would criticize in the morning, criticize in the afternoon, criticize in the evening. That was his personality.

He drew this idea to his logical conclusions by saying that (Marx (1875))

after the productive forces have also increased with the all-around development of the individual, and all the springs of co-operative wealth flow more abundantly – only then can the narrow horizon of bourgeois right be crossed in its entirety and society inscribe on its banners: From each according to his ability, to each according to his needs!

Marx was consistent during his entire life with the idealistic statement that if you change the structure of society you also change man; that is, consistent in his claim that human nature does not exist.

1.4. Three fundamental changes. A biological foundation of economics will have to produce three essential changes with respect to the existing methods used by economists. They are given by neuroeconomics, genetics and personality theory. All three are going to be described below. I will first give a brief overview of their role in economic analysis.

Let's start with the third in the list. Personality theory analyzes factors or traits explaining behavior. These factors are individual characteristics, relatively persistent over time and across situations, that are important to describe and predict behavior. These characteristics are related to what in game theory are called types of players. However, whereas a type describes the characteristics of an individual that are relevant for the specific game under consideration, personality theory identifies traits that are relevant across all strategic and economic situations. For example, the intelligence of a person is important to define the type of an individual in almost all possible strategic situations. Intelligence operates through multiple channels; it affects his payoff (for instance it has typically large effects on his income), but also affects the way in which he thinks of the situation, and the way in which he learns from past experience.

There are many other similarly important characteristics. As economists we have already considered some of them. There are others that are coming to us from psychology. Thanks to the developments we mentioned, there is now a scientific way to approach the issue of types, which so far has been studied as an abstract question. We have to start thinking of man as having a nature and thinking systematically about the fact that my nature is different from the nature of someone else. By nature, I mean our characteristics and the foundation of that characteristic here is genetic at least in part. The way in which genetics describes personalities is coming through neuroeconomics. These three themes, although very different, they are all interrelated.

Neuroeconomics is a discipline that studies the neural basis of economic and strategic behavior. It starts with the observation that in whatever decision we take, the choice made has to be the outcome of a process in our head,

after a neural process has taken place. We have begun to understand how the process operates. Neuroeconomics is usually identified with scanning the brain and see what happens in the brain while you are taking decisions. That is part of neuroeconomics, but neuroeconomics is much more than that. Research in the field is performed with functional magnetic resonance imaging (fMRI) but also done from recording from single neurons, EEG and MEG. There are many tools that one can use, but the tools are less important than the concept that we want to understand: the process that produces choices. And there are other components, for instance hormonal ones.

1.5. How these transformations affect economic analysis. Let's consider a list of examples showing how our considering the biological foundation of behavior has changed our understanding of important issues in economic analysis.

The most revealing example is our current understanding of the neural implementation of the value function, used in the analysis of multi-period maximization problems. These problems are pervasive in economic analysis. Abstractly, they can be described as follows. A state variable affects in every period the current utility. In every period we can choose an action that also affects that current utility, but also affects the probability of the transition to a new state. Thus, when we take an action today and want to fully take into account the consequences on our future utility, we have to balance the effect on current utility and future state (and thus indirectly on future utilities).

In economic analysis we reduce a potentially very complex inter-temporal choice problem using the Bellman equation, that describes how the future total discounted utility depends on the current state. The Bellman equation states that the value today of a state is the sum of the utility today, plus the discounted continuation value, at the optimal choice.

The Bellman equation is an abstract mathematical concept. To the surprise of many theorists, in the 90's experimental findings from single neuron recording in monkeys were reported, showing that learning the value function is implemented in the brain by a neural process. There are formal models, in particular Q -learning (Watkins and Dayan (1992)) models, that make the learning model formal and precise. The neural evidence shows that learning in the brain takes the form of Q -learning. We report on this in section (1.7) below.

The second example is the neural basis of choice between options. Consider an individual who is presented with a choice between two options. Economic analysis models this situation as follows: the individual computes the utilities of the two options and chooses the one which results in the maximum utility. But even this process is much more complicated than it sounds, if one wants to explain how this process takes place in the brain. Now we have evidence of how this computation is done in the brain. The

analysis is even more interesting in the case in which the experimenter does not describe the options to the subject, but instead offers a choice between two lotteries, without specifying the probability of having a specific outcome in these lotteries, or the values of the outcomes, or both.

This choice is implemented by a network of neurons organized in sub-networks, each sub-network fulfilling a specific function. Such complex interaction among sub-networks finally produces the choice among options. We can use this model to compare predictions and observations made both in choice data and neural data. We will examine this research closely in the section on choice below. The formal models are based on the theory of stochastic diffusion processes. As mathematical models they are not very complex, but they become a lot more interesting when you appreciate that they approximate the process that takes place in the brain when you make a choice.

A third example is the adaptation to the environment. We illustrate first this idea with an example that will be familiar to everyone. When someone goes from a brightly lit environment into a very dark room, for a few seconds perception is greatly reduced. Since the intensity of light is now much weaker, retina and brain processes have to adjust to the new situation. The fact that this adjustment happens in visual perception is everyday experience. Something similar happens in choice and in reward processing. We are now beginning to understand how this adjustment works.

1.6. Biases and their practical implications. The similarity of what happens in visual perception with what occurs in reward evaluation suggests a general consideration regarding the policy implications of the research in these topics, and in particular those that typically derived from research in behavioral economics.

Illusions in visual perception are ubiquitous. All of us have observed since infancy that if one puts a straw in a glass of water, and looks at the glass, the stick appears to be broken. This simple fact does not make society think that there should be special interventions to help people against the consequences of such visual illusion. We do not have a Ministry of Correct Visual Perception. The reason is clear: even if there are distortions in visual perception, we are usually very effective in adjusting to it. Even when you show me a straw that appears broken, and I conclude since the age of three that I cannot trust my visual perception blindly, I still go out in the street and I'm pretty confident that I will not be run over by a bus.

Thus, to evaluate the practical policy importance of a bias (in visual perception as well as in economic behavior) it is important to estimate two very precise quantities: how large the effect is and how long it lasts. This estimation is usually ignored when researchers report the existence of biases; apparently the simple existence of a bias is sufficient to show the need for a policy correction.

We can now consider an example that shows how, for matters of fundamental importance, our neural systems function amazingly well.

1.7. Value function learning. The example is by now an important piece in the history of neuroeconomics (Schultz et al. (1997)). A researcher may record neuronal spikes in a monkey's brain in response to rewards. A monkey is observing a screen, and then, unexpectedly, receives a reward. The delivery of the reward induces a sudden burst of spikes in reward-related neurons, which can be recorded. Note that in this case when there is no conditional stimulus announcing the coming reward; the delivery of reward was unexpected. Now consider a different experiment: 10 seconds before the delivery of the reward, a researcher displays on the screen to that same monkey a stimulus, a sign in itself meaningless; let's say a triangle. He then exposes the monkey repeatedly to this sequence of presentation of stimulus and delivery of reward. Since there is now a systematic association between the otherwise meaningless symbol and the delivery of a reward, slowly the monkey learns that 10 seconds after it sees the triangle it will receive a reward. The expectation of the reward is now adjusted to the point in which this information is acquired, firing up the neurons.

Let us now go back to Q -learning. In this model learning the value function is a process which at each period gives the new value as a linear combination of the value function of the previous period, plus an adjustment term (called prediction error) equal to the difference between realized and expected (according to the old value function) value. The neural response we have described performs exactly such prediction error adjustment. Q -learning is what the brain is doing.

One can run further tests with variations of this experiment: for example one can give a conditioned stimulus but offer the reward 10 seconds later than the "promised" time. Now the neural spikes at that time are reduced because the adjustment is negative, since an announced reward was not delivered. In this case too the neural response is coding the prediction error.

1.8. Regret and envy. Let's see how this view of the learning process can help us explain emotions. We will analyze this topic here in some detail because we will not consider it more extensively later.

We usually express the emotion of regret with a statement like "I wish I had not done something in the past". When one views it in this light, regret appears to be an irrational, ineffective emotion, because it attaches emotional significance to events that are by now in the past and thus impossible to modify. Since emotional processes are expensive, for the energy they require and the consequences they induce, it is natural to ask why such emotions should arise, and how it could persist.

A different way to look at this emotion is suggested instead by considering it for what it is, namely counterfactual thinking. From this point of view, when you have a decision, you first select a choice, then you observe the

outcome, and then you go back and perform the counterfactual estimation: was it a good idea to choose what I did? If it happens once in a lifetime, regret is really useless, but if you have several opportunities to make the choice, regret is very useful since it's teaching you on the basis of the outcome you realize whether the action you took was good or not. Regret can then be considered a step in learning a value function.

Now think of a situation in which you made a choice, one of your friends made the other choice, and he performed better than you. In this case, you are likely to experience an emotion of envy. The similarity with regret should be clear. You had a decision to make, you made a choice, and you evaluate the consequences. The fact that someone else had the same decision to make, took a different option, and had a different outcome makes the counterfactual evaluation more natural, almost inevitable, because the counterfactual outcome is readily available to you, even if you do not go through the thinking process, or the factual investigation, of discovering what the consequence of different choice would have been.

If we analyze envy from this point of view, it seems clear that we may consider envy is the social correspondent of regret, and both as an instance of learning, of which we can provide the biological basis. As you learn counter-factually that you could have done something better, you can learn counter-factually that you could have done something better following the example of another guy. Envy is not a particularly good emotion, but we have it and it has not been selected away because it fulfills this learning function in social environments.

This point of view provides predictions on brain activity that is associated with envy and regret. The first prediction is that both emotions should involve neural structures that are associated with the evaluation of rewards. If the role of regret and envy is similar, then the neural activation in the two cases should be similar. Some further thought may suggest that envy might have additional components, which are also easy to derive from the functional role of these emotions. Envy should code, in addition to the counterfactual evaluation of one's past actions, also the effect on the current social status of a person.

These predictions can be tested, and when they are the results provide support for the predictions. My co-authors and I (see Coricelli and Rustichini (2010), Bault et al. (2011)) gathered experimental data on choices, verbal assessment of subjective states, and brain activity in subjects that were choosing lotteries in a *private condition* (in which they observed the outcome of the chosen and of the unchosen lottery, a condition likely to induce regret in the case of a bad outcome), and in a *social condition* (in which they observed the outcome of the lottery chosen by another person, a condition likely instead to induce envy in the case of a bad outcome.)

Confirming the first prediction, the striatum, a reward-related brain structure, showed activation in both conditions (private and social) at the stage of observing outcomes. Confirming the second prediction, we observed higher

activity when participants won more than their counterpart (thus having social, rather than just private, gains) compared with winning in isolation. They also had lower activity when they won less than their counterpart (a social loss), compared to a private loss.

1.9. Summary so far. We have reached two conclusions. First, one can replace “as if” models with mechanistic models of choice. Friedman in his 1953 book “Essays in positive economics” (Friedman (1953)) argued it doesn’t matter whether the model is correct or not as long it gives good prediction (thus, “as if” it were correct). This philosophical attitude is profoundly wrong. It is important to understand how things work because eventually we will be able to make prediction of the sample only if we understand how the process works.

Second (and this is philosophically very important): economists may be inclined to think that our choices are optimal. Instead, it is much more plausible that the process producing choice, not the choices, are optimal. The choice process is the only entity that evolution and selection can control and modify. Evolution and selection cannot intervene and decide for you whether to choose this option or do this action (for example, whether to choose this college degree or the other). It can only teach you what how to cope best within a constantly changing environment.

1.10. Genoeconomics. Genoeconomics is the discipline that studies genetic basis of economic and strategic behavior. It has become a systematic research program after the year 2000 when, with the Human Genome Project, a complete description of the human genome has become a reality.

Let us first recall the basic notions. Our genome (the complete set of genetic instructions) is inherited from two parents. In humans, these instructions are organized in two sets of pairwise similar chromosomes. Each chromosome consists of two strands. Each strand in turn results of sequences made of four nucleotides, labeled A, T, G and C. A in one strand is always paired with T and G with C. Therefore, if I give you the sequence of nucleotides on one strand you can reconstruct that on the other yourself — wherever you find the T you know that in the other is an A and so on. Thus one of the strands is completely irrelevant so you can ignore it.

Only in about 1 out of 1000 loci there is a variability of nucleotide in the locus. For instance, there may be a variation between the strand that is coming from your father and the one that is coming from your mother, there may be a T and a C in a position. This variability is a polymorphism, which can be translated as “multiple forms” from the original Greek.

1.11. Genotype for each subject-GWA Studies. Data on genotype of an individual list, for each subject, chromosome number, position in a gene that we are considering, and the corresponding number of A, G, T and C nucleotides in that position. Since in each position we may have two of the four nucleotides, the coding is very long and complex. It is partially

simplified by the fact we mentioned that in each position, typically, only at most two of the four nucleotides appear in the population.

If there is only one (for example, if in all individuals of a species there is only a pair (A,A) in a position), then there is no variability in that position and we do not need to track the individual values there. Consider now the interesting possible case which is left, in which there are two possible, say T and C. So in that position we may have pairs (T,T), (T,C) or (C,C), each occurring with non-negligible probability. In this case there is, in the species we are considering, a variability or a polymorphism at that position. Since this variability takes the form of two single nucleotides that may be found, we call this variability single nucleotide polymorphism (*SNP*).

An equivalent coding, based on the knowledge that at that position we have either T or C, is to keep track of the number of T's at that position. We can adopt the convention that if you have one T and one C, we assign a value 0.5 to that position, if you have two C's we assign a 0, and a 1 if you have two T's.

Now if for every individual we have information on, say, whether you had college or not, we can run a uni-variate regression of the phenotype of interest on the value of the SNP at that position on millions of data. For instance let's take a position and a dependent variable which is, say, whether this subject had college. We can run a logistic regression of whether the individual has college or not on this independent value. This estimate will give us a coefficient which is our estimate of the effect of this SNP on the probability of getting college. In general such association studies regress a dependent variable of interest on a very long list of variables, each describing the genetic profile of individuals at each position, and producing the GWAS coefficient for that characteristic or phenotype.

1.12. Polygenic scores. If we consider a given phenotype, a polygenic score (*PGS*) is a single real number, individual specific, that describes the genetic value for that phenotype in an individual with that score. If the phenotype is a disease, then the polygenic score is usually called a polygenic *risk* score, and it gives a measure of the disease risk due to genes of an individual. The polygenic score can be constructed as follows.

We have a very long list of values of the SNP's, one list for each subject. For each of these SNP's, we can estimate a coefficient telling us how likely a person is (for example) to achieve college education depending on whether the person has a value of 0, $\frac{1}{2}$ or 1 in that position. The coefficients may be estimated from a logistic (if the phenotype has only two values, as for college) or an ordinary linear regression. From this estimate we derive a very long list of coefficients, which hold for the entire population. This is a Genome Wide Association Study (GWAS).

For each subject in a sample of interest, which is not included in the sample used to estimate the coefficients, we again have a sequence of 0, $\frac{1}{2}$ or 1 for each locus. We can compute the polygenic score by taking the inner

product (that is the sum of the product at each coordinate) of the vector of coefficients and *SNP* values of each individual. A correct computation of the score has to consider some additional features, and perform corrections due for example to the correlation of the *SNP* values at different position. But the general idea is the one described here.

1.13. PGS and social mobility. In (Rustichini et al. (2023)), we applied these insights to analyze causes of income and education, and mobility in income and education with respect to parents, using together with other information the polygenic score of individuals. We take as starting point a very standard model of intergenerational mobility like (Becker and Tomes (1979)) and test it using the Minnesota Twin Family Study (MTFS) data. These data are described for instance in Iacono et al. (1999) or Disney et al. (1999). In the model we consider, parents are altruistic to their children, and decide how much they invest in their education depending on the available income. Thus, the human capital of twin $j = 1, 2$ in family i is given by the parental investment, an error term and a term θ which is the skill of the individual.

We introduce in this setup a substantial change: the variable θ does not follow an $AR(1)$ process, but a process determined by the rules of inheritance derived from fully specified genetic model. In addition, while in standard models derived in the tradition of (Becker and Tomes (1979)) this skill is a latent variables, we have a measurement of it provided by the polygenic score for EA for the individual.

This component of the model describes how this term θ evolves over time. We also can infer the properties of this process for members of the family. The θ of a person is obviously related to that of the parents, for whom we also have the polygenic score. Also, if we have two twins, and the twins are identical, then the θ is going to be the same. If the are not identical twins, there should be a correlation of one half between the two scores. Once we specify θ to be described by the polygenic score we have a substantially richer model.

We can use the laws of genetic transmission to define the stochastic process governing the time evolution of θ ; but we must first specify how the parents are matched. For example, an intelligent mother, or a mother with a higher polygenic EA score, is more likely (and we can estimate exactly how much more likely) matched with a more intelligent father, or higher polygenic score. Therefore, the process of combining the two genotypes of mother and father is very different from an $AR(1)$ process. In the MTFS data we have information on the genotype of the parents and their twin children, so we can study this process on the basis of the knowledge of all the variables.

To complete the model we need to decide what is the stable matching, that is how the pairing of parents results in a long run equilibrium. This is usually the result of a combination of the income and the genotype of the

two. For instance, it is well known that the intelligence of the parents is highly correlated. We then characterize the resulting non-linear stochastic process and study its invariant measure.

Some properties of the invariant measure are interesting and illuminating. For example we can plot frequency in the population of different alleles against the income. We can then consider the income of a person and, for each person, the frequency of the SNP favorable to educational attainment and compare it with the frequency of an irrelevant SNP. The frequency of the allele that is enhancing educational attainment increases with income, and the steepness of the curve is larger, the larger the coefficient of that allele in the GWAS for educational attainment.

I described here a very complex but realistic process in which parents meet depending on their intelligence and other important traits on the one hand, and their income on the other. The matching is assortative; this complicated process tends to twist the invariant measure away from random association and it will make the correlation between income and different SNPs adjust. The adjustment is stronger for SNPs which have an effect on educational attainment.

1.14. The role of the polygenic score. Regression analysis shows that the polygenic score matters for education, as should be expected, but also for income. Consider for example table (1) below.

Model (1) of the table reports the values after control for sex and interaction between sex and family income. Age has not a significant effect, likely because age in the sample is very similar across subjects. Sex has a strong and significant effect: income for male individuals has a substantially larger intercept (27.7 per cent), but a smaller (by 6 per cent) dependence on the family income.

In Model (2), the coefficient of the individual polygenic score is 7.8 per cent (SE = 0.025, p -value = 0.002). Its size is approximately half of that of family income (12.8 per cent). Considering that the polygenic score we are using is estimated from coefficients from a GWAS for education, it is likely that the weight of genetic factors affecting income is higher.

Finally, Model (3) in the table presents controls for some of the variables that are likely to mediate the effect of the polygenic score. Education Years is the most natural variable to measure the effect of the polygenic score in education, and in fact the estimated coefficient is large (25.6 per cent, (SE = 0.035), p -value < 0.001) and significant.

It is important to understand precisely what are the causal pathways from the genotype of the individual, as summarized by the polygenic score. To do this we can construct a structural equations model (*SEM*) describing the role of potential mediators of the effect. The structural component of such *SEM* estimation is reported in table 2.

In the equation for education years, the coefficient both C and NC are significant. We can compute with the delta method the product of the

coefficient for the link from the PGS to the variable C , times the coefficient from C to Education Years. The value of the product is 0.082, (SE = 0.018, $z = 4.53$, p -value < 0.001), with confidence interval [0.046, 0.117]. The corresponding product for the path passing through NC has a value of 0.034, (SE = 0.019, $z = 1.8$, p -value = 0.071), with confidence interval [-0.003, 0.072].

Once we control for C and NC , the coefficient of the PGS is not significant (p -value = 0.725). For comparison we note that in the regression restricted to twins, controlling only for sex, the coefficient is 18.7 per cent (SE = 0.022, $z = 8.37$, p -value < 0.001). The coefficients of education of

TABLE 1. Income at the age 29 years of age, on family income, PGS, and Personality. The dependent variable is a measure of income at 29. All variables, including college of parents and male, are standardized to mean zero and SD 1. The signs of MPQ (Multidimensional Personality Questionnaire) variables NA (Negative Affect), Externalizing and Academic problems are reversed. Controlled for PC's (principal components) and parents-child time difference in age at income data collection.

	(1) b/se	(2) b/se	(3) b/se
Family Income	0.134*** (0.027)	0.128*** (0.027)	0.078** (0.032)
Male	0.277*** (0.025)	0.276*** (0.025)	0.313*** (0.029)
Male \times Family Income	-0.060** (0.025)	-0.060** (0.025)	-0.050* (0.030)
PGS		0.078*** (0.025)	0.021 (0.028)
Education Years			0.256*** (0.035)
IQ			0.008 (0.029)
MPQ PA			0.061** (0.026)
MPQ NA			-0.024 (0.027)
MPQ CN			0.034 (0.032)
Externalizing			-0.072* (0.037)
Academic effort			0.057 (0.038)
Academic problems			-0.017 (0.034)
N	2100	2100	1485

parents and family income are both significant, of the same order of magnitude, but education of parents (13.6 per cent (SE = 0.029, $z = 4.58$, p -value < 0.001)) is approximately twice that of family income (7.5 per cent (SE = 0.031, $z = 2.38$, p -value = 0.017)). The *PGS* of parents is not significant.

1.15. Personality theory. We now come to the third of the ingredients in the transformation of economics into a discipline with a biological foundation: the conceptual organization of individual types into personality traits. This conceptual structure makes the analysis of human nature concrete, specific, and quantitatively precise.

No one today questions the idea that all men are equal, if we understand the statement to mean that all men have equal rights. But men are not equal in their personal characteristics: human beings are very different in skills, preferences and general attitudes. Psychology has done great progress in the last decades in building a theory of personality, that is of characteristics of

TABLE 2. *SEM of Pathways from PGS to Education Years.* See (Rustichini et al. (2023)) for a description in detail of the *SEM* model estimated, equation (37) to (41). All observed variables standardized to mean zero and SD 1. Cognitive skills test scores (*ct*'s) are verbal and performance IQ, non-cognitive (*nct*'s) are the three broad *MPQ* dimensions. Standard errors estimated by bootstrapping. $N = 852$. Model *vs* saturated: $Pr > \chi^2 < 0.0001$.

Equation	Variable	b	z	p value	CI
Ed Yrs	C	0.285 (0.058)	4.87	<0.001	[0.171, 0.401]
	NC	0.856 (0.276)	3.11	0.002	[0.315 , 1.4397]
	PGS	0.014 (0.041)	0.35	0.725	[-0.066 , 0.94]
	PGS mother	0.033 (0.030)	0.71	0.282	[-0.027 , 0.093]
	PGS father	0.019 (0.030)	0.66	0.512	[-0.039 , 0.078]
	Educ Parents	0.136 (0.29)	4.58	<0.001	[0.078 , 0.194]
	Family Income	0.075 (0.031)	2.38	0.017	[0.013 , 0.137]
	Male	-0.151 (0.055)	-2.77	0.007	[-0.260 , -0.041]
	Constant	0.376 (0.027)	9.85	<0.001	[0.301 , 0.450]
	C	PGS	0.287 (0.031)	9.21	<0.001
NC	PGS	0.040 (0.025)	1.95	0.051	[-0.0002, 0.081]

individual that are stable over time and situations, and affect significantly the individual's behavior.

In economic analysis we have for a long time typically used two variables that describe the attitude of a person: the attitude to risk, relevant for choices with uncertain outcomes, and the discount factor, relevant for choices over option delivering rewards at different points in time. But there are many other characteristics that affect human and also economic behavior. Some of them with a prominent role in public discussion, such as intelligence. Another one is conscientiousness, that is how systematic you are in the pursuit of your targets; you are conscientious if the night before the exam you don't go to a party.

These theories can be tested and used for interpretation. We have information on personality traits of each individual from several sources, and this information can be used to predict their behavior in a variety of situations.

We need to understand the role of personality in behavior, and then we can link personality theory with neuroeconomics and genetics to understand the causal structure of important phenomena, for instance the deep reasons for the existence of inequality and social mobility. In doing this, an investigation of the role of different skills in economic success of individuals is essential.

We now consider some examples of research in this direction.

1.16. College education and intelligence. In a recent paper (Ichino et al. (2022)) my co-authors and I study the profile of intelligence in the UK universities from 1960 to 2000. The question is important because, as the access to the university in England between 1960 and 2000 has become easier, the level of intelligence of people in college on average is likely to have gone down. This conclusion may appear natural but it is far from obvious. In fact such conclusion is in sharp contrast with the analysis of Robbins in his report (Committee on Higher Education; Great Britain (1963)) which provided the justification for policies of expansion of the university in the 1960's. The Robbins report was based on the idea that it was possible to expand access to University and at the same time be able to pick very capable students coming from the less privileged sectors of society (a wealth of "untapped talent"). As a result, the level of intelligence in the University would not necessarily decline with the expansion, or might even go up. Our results in (Ichino et al. (2022)) show a very different conclusion, and thus should prompt us to be careful when we evaluate projects of large expansions in higher education.

1.17. Behavior in games. My coauthors and I have experimentally studied in a set of of papers the behavior of subjects in repeated games. (Proto et al. (2019), Proto et al. (2022)).

In these experiments you ask subjects to play a game like the prisoner dilemma repeatedly, with parameters of the experiments such that the repeated game has multiple equilibria: for instance "defect defect" repeatedly is an equilibrium, but "cooperate cooperate" is also an equilibrium.

In the design we adopted, in the first day of the experiment you bring the subjects to the laboratory and ask them a set of questions. Some of them give a measure of intelligence. In a second day, now that we know how intelligent each person is, we run a repeated prisoner dilemma experimentally, but we allocate the more intelligent and the less intelligent in separate groups, so that we can observe the role of different intelligence in strategic behavior. If we set the parameters in such a way that the only equilibrium is “defect-defect”, then we observe intelligent people defect strongly; thus they do not have any unconditional inclination to cooperate. The more interesting question is what happens in range of parameters where you can either always defect or always cooperate.

Our main finding in (Proto et al. (2019)) is that intelligent people cooperate more than less intelligent people, when gains from cooperation are possible. The link is not mechanical, and intelligent people are not unconditional cooperators. For example, when the discount factor is low, so that cooperation is not an equilibrium, they choose defection in the repeated Prisoner’s Dilemma just as subjects with lower intelligence. In (Proto et al. (2022)) we find that intelligence operates by reducing the frequency of errors made in implementing a repeated game strategy: higher intelligence subjects exhibit lower frequency of errors.

We conclude that important regularities in strategic behavior, some of which cannot be explained by game theory (because that theory is sometimes unable to uniquely predict behavior in repeated games, as in the prisoner’s dilemma with high continuation probability), can instead be explained considering personality theory and information about characteristics of the individual.

1.18. Political behavior. The next question we discuss is whether genetic analysis can help to explain political behavior; for example, participation in elections or party affiliation. There is a growing research on this topic. I will illustrate the nature of insights one can derive by considering how genetic characteristics, that can explain at least in part the educational attainment of an individual, can also contribute to explain the political behavior. This is an interesting application because it shows how we can combine information from personality theory (and in particular the role of intelligence or other factors contributing to success in education) and genetics.

In a recent paper (Dawes et al. (2021)) we consider data from the US, UK and Sweden reporting whether a person votes in election, and in particular whether he votes in important or less important elections. For example, in the case of a USA citizen, we take into account whether he votes in presidential versus midterm elections. We also have available estimates of the polygenic score for educational attainment (EA).

We investigate this question: if a person has a higher EA polygenic score than another, which of the two is more likely to vote, and in which type of

elections? We found that intelligent people with higher polygenic score vote more and they vote more in the less relevant elections.

The explanation of this finding seems natural: voting is information processing, which is costly. Thus, you are less likely to vote the more costly this processing is. The statement that voting is information processing seems natural, and one can give formal justification of it within rational choice models (see for example Feddersen and Pesendorfer (1996), McMurray (2013)). The processing is less costly if you are smarter, hence our findings. So there are sound and profound reasons why the attitude to politics is profoundly affected by our genes, and we can begin to see the direction in which the effect operates.

1.19. Summary. I described three directions in which economics is changing, and should change even more than it has so far; and I have given some examples of what has been done. Thanks to neuro-economics and genetics we have a new foundation, the most coherent plan to implement this project.

It is now time to examine some of these ideas more in detail.

2. SELECTION AND THE ROY MODEL

In this section we illustrate how one can use of tools from genetics to understand the evolution of the personality characteristics of a population over time. Our conceptual tool, derived from economics, will be the Roy model (Roy (1951)). The model was developed to predict how people with different skills will choose one of two possible activities (hunting or fishing). The paper finds that the allocation depends on the distribution of skills, the correlation between these skills in the population, and the productivity of the two available technologies.

In Roy (1951), the distribution of skills is given and exogenous, and the model ignores the reasons of the skill differences among individuals. We add to the setup of the Roy model the biological component (Jandarova and Rustichini (2023)). In this case, the crucial fact is that the distribution of skills and their correlation is determined by the distribution of alleles in the population; in particular, it is endogenous and variable. If there is a change in the productivity of the technologies, then the distribution of skills is no longer exogenous and given but will change over time because the relative productivity in the two sectors will affect the reproductive fitness of people with different genotype.

The specific hypothesis we test is the following. In the recent past, starting from a period between 14000 to 12000 years before present (BP), there has been a selective force operating in the direction of an increase of the frequency of educational attainment enhancing alleles. Selection has operated in the last 14000 years by changing the distribution of the genotype among humans, as human societies transitioned from foraging hunter-gatherer society to farmers during the Neolithic Agricultural Transition (NAT).

This is an instance of a more general hypothesis, that technical change may induce a genetic change in the characteristics of a population, through the modification of the frequency of alleles in the population. Perhaps the initial technical change is induced by a change in relative productivity of available technologies. This is the case in our example of NAT, which was probably due to climate change. But later further changes in technology are compounded with the genetic changes first and institutional changes second. In this view of history, technical change induces a genetic change, which in turn induces a further change in economic activity and institutions.

2.1. Technology, genetics and institutions. What we suggest here is a special instance of a general relation that we hypothesize between technology, genetics and institutions. Technical change may operate in many different ways. Change in medical technology, for instance, may alter the natality and mortality rates possible in a society with a given level of development of productive forces. In this case the pathway of change may operate not through a Roy type change in allocation of activity but through mass migration.

A comparison with a different broad view of human history may be useful. Marx proposed that technical progress and the development of material forces of production changes the relations of production or even the mode of production. These transformations would then induce a change in the institutions of a society. He ignored the role of genetics, of course. The first reason is, of course, that he did not know about it. The second is that, as we stress later, he was an idealist philosopher, the last of the Hegelian left. This omission may be forgivable in his case, at least because the understanding of these factors was so limited; it is no longer excusable today.

Let us see in detail how these transformations occurred in our distant past.

2.2. Neolithic agricultural transition (NAT). The Neolithic agricultural transition started about 12-11,000 years ago. It was a transition from a lifestyle of hunting and gathering to one of agriculture and settlement. There was an earlier transformation from foraging into farming among Natufians, probably also induced by an earlier and more modest climate change. This transformation involved a widespread process of genetic transformation of plants and animals. This process is called domestication, as for example is the change of wolves into dogs. Everybody will agree that the transformation from wolf to dog was not cultural: in fact we all agree a dog is a different animal. It is, therefore, reasonable to conjecture that a similar process occurred in human population. This transformation originated in the Fertile Crescent and spread to Europe at the speed of one kilometer per year.

2.3. Neolithic demographic transition (NDT). The Neolithic agricultural transformation was associated with a demographic transition. Such

demographic transition can be detected (Bocquet-Appel and Masset (1982), Bocquet-Appel and Bar-Yosef (2008), Bocquet-Appel (2011), Bocquet-Appel and Degioanni (2013)) with the study of the proportion of immature skeletons (age between 5 and 15 years) in cemeteries. This fraction showed a sudden increase around that period, as society changed from nomadic to sedentary. This increase can be explained well by a model in which there is a sudden increase in fertility which was in turn induced by a change in the energy balance of the mothers. Earlier the mothers were consuming less because they were sedentary, and when their diet changed from those that of foragers to that of farming, with higher fraction of carbohydrates. So both intake and outtake produced an increase in fertility of women, causing increase in population, which was then compensated by an increase in mortality induced by crowded living conditions. This is the Malthus model in action.

It is a time sequence which is opposite to the contemporary demographic transition, which in the last 200 years saw a decline in mortality, an increase in population, and finally a decline in fertility. Population size is now becoming stable and will be declining by about the middle of this century.

2.4. A different view: NAT and NDT as institutional changes.

Economists have considered in the past the problem of the agricultural and demographic transition. The way in which they have analyzed them is a good illustration of how economics without a biological foundation may completely miss not just the explanation, but the question itself. Let us see how they proceed.

They start (see Bowles (2011), Bowles and Choi (2013), Rowthorn (2011), Robson (2010)) from a puzzle which we briefly recall. Some archaeological evidence based on study skeletons of foragers and farmers, suggests that the health condition of foragers was not substantially worse than that of farmers; in fact in some respects it was even better. Also the productivity of agricultural products among foragers was not much lower than among farmers (Bowles (2011)). The reality and the interpretation of this evidence is debated, but economists focused on these findings, and I would like to take this as an example of how economists end up explaining an event when they choose to ignore the biological side.

The main claim made in this literature is that the key transformation within this historical event was the emergence of private property. The thesis proposed is the following: when you change from foraging to farming you plant seeds, and you raise herd of cattle. So now you have to guarantee that the output of the seeds you plant and the animals you raise is yours; you have to protect your produce from theft. They use this fact to explain the following apparent puzzle: if the claim of inferior productivity and equal if not worse health conditions is true, then society seems to have made an inferior choice going from foraging technology to farming technology.

Rowthorn and Seabright (2010) develop formally this explanation in a clear and simple way, and I will summarize their reasoning here.

Define a simple symmetric two-players, two-actions game in which two groups, each acting collectively as a player in the game, can choose either foraging or farming. Let us assume, consistently with the claims recalled earlier, that both players choosing foraging gives larger utility to both than both choosing farming. In this game, foraging is equivalent to cooperating and farming to defecting. If the product from both players foraging is better for both players than what they get when both players choose farming, then how do we end up in an equilibrium where both players farm? The underlying logic of the answer is the same operating in the prisoner's dilemma. We can indeed verify that the game we have described is a prisoner's dilemma, by checking what happens if you stick to foraging and I switch to farming. The prisoner's dilemma logic requires my payoff to be higher as a consequence of the switch, and indeed it is. The reason is simple. If the other player insists in foraging and you go from foraging to farming, you must build a military structure that is necessary for the defense of the private property, your seeds and your cattle, and this is a cost. But once you do that, the same military structure that can be used for defensive purposes can now be used for offensive purposes: you can now easily steal goods from the other. All things considered, if the other is still foraging, I am better off choosing farming. This is the explanation of the puzzle.

The exam of the institutional changes which are claimed to underlie the agricultural and demographic neolithic transitions suggests two points that illustrate our main thesis, namely that we need an economic science with a biological foundation.

2.5. The methodology of focused investigation. The first is methodological. There is a famous little story about a fictional character, Mullah Nasreddin, struggling to find his lost ring. For a while he searched for it in the living room, where he thought he had lost it, but he could not find it. He then went out in the yard and began to look there. His wife said "Mullah, you lost your ring in the living room, why are you looking for it in the yard?" Mullah answered "The room is too dark and I can't see very well. I came out in the courtyard to look for my ring because there is much more light."

Economists are victims of the same indulgence. They know how to study institutional change, so they focus their investigation on institutional change, even if the most likely explanation of the events is a biological one. In the case we are considering, the change is mediated by a genetic transformation. But a genetic transformation is a very complex phenomenon, which requires vast technical knowledge. Genetics, in our story, a room which is very dark. Although the ring is there, that room is very dark, and so economists look elsewhere for an explanation. This is a methodological shortcoming of economics that we might want to address.

2.6. The ideology. The second point is related to the political philosophies behind the different views of economics today. The following passage from Rowthorne and Seabright (2010) is indicative of how sometimes economists fall into a well known trap. Discussing the evidence that early neolithic living standards (in which people were surviving though foraging) and comparing this condition with what followed after the agricultural revolution, the author long for a lost state of nature:

“The evidence about early Neolithic living standards (that is the evidence that foraging was better than farming) add substance to the eternal appeal that myths of the noble savage have had throughout human history, since such myths have seemed to suggest, counter-intuitively, that economic development since the time of the alleged fall has been both inevitable and regrettable.”

The Rousseauian inspiration of this passage is hard to miss; in particular the evocation of the myth of the noble savage in the second discourse in inequality (Rousseau (1999)). Evaluating a transformation (from foraging to farming) that has allowed human civilization to evolve out of a condition of permanent hunger, two economists linger on the romantic past.

2.7. Alternative interpretation and modeling of two transitions.

Let us now consider a possible different explanation (Jandarova and Rustichini (2023)). We start with a fact on climate A climate change occurred about twelve thousand years ago; this fact is well documented. The change of climate altered the relative productivity of foraging and farming, and thus induced a change in the allocation of individuals of different skills to different occupations. If the temperature rises, then the productivity of farming changes (likely, it increases) with respect to foraging. The difference in productivity induces a difference in reproductive fitness of farming community, so the relative number of farmers increases. We have seen from the analysis of the Neolithic demographic transition that there was also, in that period, a demographic change; this change in turn produced a change in the distribution of the genotype. This created larger concentrated communities, cities and new institutions.

That is the instance of a process I had described earlier: a change in technology inducing a change in genetic composition and then changing institutions. This is a hypothesis we may test using genetic data. Let us now consider the different components of this scenario.

2.8. Climate change. The appearance of a flower with eight petals, the dryas (*Dryas octopetala*), is reliable evidence that the temperature has fallen. This occurred in the period between 12900 and 11600 years ago (thus called the Younger Dryas). This was the last period of low temperature. When considering temperature change, researchers commonly report data on the temperature of the ice core of Greenland. Greenland is far from the Fertile

Crescent, so the fact the temperature changed in one region might not be sufficient proof that it also changed in the other. But there are data that we report which show that the change occurred even in the fertile crescent, and it was large. For example, the Mediterranean sea surface temperature in Celsius went from around 10 to 23. Such a huge climate change of course changed the relative productivity of agriculture and foraging. Other evidence, such as the temperature in the Soreq caves in Israel show evidence of an increase. (see Robinson et al. (2006) for a survey).

2.9. The estimation problem. Now we can consider our estimation problem and hypothesis testing. We use the same structure as in Rustichini et al. (2023), adapting it to the present task. The data we use are ancient DNA (aDNA) data, that is DNA isolated from ancient human specimens. A model will provide us with a stochastic process of the distribution of alleles in a population subject to selection. The characteristics of the stochastic process depend on parameters: the matching of predicted and observed data allows us to estimate parameters, and their standard errors.

Our estimation problem is the following. We take the frequency of alleles that have been discovered as relevant for educational attainment in modern population. We take in hunter-gatherer population as our initial condition and the contemporary frequency as the final condition. We provide a selection model that produces predicted data. A parameter to be estimated measures the direction of the strength of selection. We compute the distance between the predicted final frequency according to the model and the contemporary frequency, which is the final condition. We then test the hypothesis that the strength of selection parameter is positive and significant. For simplicity we adopt the assumption that mating was random (that is, a man mates with a woman randomly, independent of the characteristic of the two individuals).

2.10. Genotypes and allele frequencies. Let me introduce some minimal notation: K is the number of loci that are relevant, that is loci in which the polymorphisms affect a phenotype. The number of individuals is M , so $2M$ is the number of total alleles. Selection in a population depends on many factors, including how large the population is. Recall that we have two strands in our genotype. If you keep the information how a SNP is associated to the other along a strand then the relevant information is the haplotype. If we ignore, then it is the genotype. If we use haplotypes, then we are considering a probability distribution on the haplotypes induced by a population of size M .

The phenotype of interest, for example intelligence or conscientiousness, or a combination of the two, as it is likely to be the case for the relevant phenotype for education attainment, is denoted by Z . It is in part explained by the polygenic score, which is obtained (as outlined in section (1.12)) as a sum of the products of GWAS coefficients and individual score at each locus.

The phenotype is eventually determined in part by the genetic component measured by this score and in part by environment.

We can then define a fitness function. In a simple model, we can take fitness to be equal to the phenotype, multiplied by a real number ω that can be positive or negative. I will use here mostly this linear model (corresponding to a directional selection). The probability according to which someone with a specific genotype is selected the next period depends on how fit the individual with that genotype is. The parameter ω is describing how in a society a phenotype is selected. If ω is positive, then a higher phenotype makes the individual more fit, and thus more likely to be selected. The crucial empirical question is whether this ω is positive or negative, and whether it changed about 12000 years ago.

2.11. GWAS coefficients. To get an idea of how these coefficients are distributed, one can plot the distribution of the effect size of the GWAS coefficients. The size of the coefficient is distributed negatively, sharper than exponential. There are about 400 SNP's that are important and a much larger number of SNP's that are less important.

Here is the first test of our model. If one plots the polygenic score for educational attainment in the population against time, one can see the following regularities. The score is approximately stable between 30000 to 10000 years in the past. At that point, it begins to increase steadily for the following years. This picture suggests some support for our hypothesis: there was a change in the distribution of the genotype which occurred around the time in which the climate changed, following the last stage of the Pleistocene epoch, called Younger Dryas (approximately 12,900 to 11,700 years BP). But this test requires that we can make comparisons of the polygenic score across different populations, which is not entirely justified, so we need a deeper analysis.

2.12. DNA data: ancient and contemporary. For the analysis we need data on the genetic profile of individuals. As we mentioned, we can use ancient DNA. A similar analysis can be done using *contemporary* DNA, rather than aDNA. This can be done relying on an application of the theory of stochastic processes, the coalescent theory (see Kingman (1982) for the methodological foundation, and Berg and Coop (2014) or Stern et al. (2021) for applications). To use contemporary DNA one can reconstruct the sequence of mutations in the past, determining the order over time in which they occurred. From the present data we can reconstruct the stochastic process in the past, going backward, and thus also estimate the polygenic score in the past. The results that emerge are largely consistent with those we get from the analysis of ancient DNA: the polygenic score was stable until 560 generations ago and started going up afterwards. The number 560 of generations is obtained by dividing the number of years (14000) by 25 years. So, everything that we want to explain must realistically happen within 560 generations.

2.13. A model of selection. The general idea of the model is the following. Consider one of the K positions mentioned in section (2.10), with two alleles, say A and C. One may count how many A's and C's there are in a population, normalize by the population size, and get a number between zero and one for the allele which has a positive coefficient, let's say A, hence relatively favorable to educational attainment. That number is the frequency of the alleles favorable to educational attainment. The entire description of the mean of the population is described by a function of time, written $p(t)$ at the beginning of a generation. This is a vector of dimension K , that is the number of relevant alleles. At the end of the sequence of events we obtain a new generation, leading to $p(t+1)$, the new frequency. This new frequency is a random variable of which we can study mean and variance. Thus, within a generation, we have an initial frequency, and then a sequence of steps to go through, which at the end produce the new frequency. The sequence of steps include the initial condition, mutations, and a Hardy-Weinberg condition (the equilibrium that you get when there is random mating). Then selection operates. Each individual has a fitness. For example, if selection favored qualities that induce higher educational attainment in the contemporary world, an individual with a higher score for this phenotype would be more likely to reproduce. This final step produces the new frequency in the next period.

This sequence of steps describes a non-linear Markov process which depends on the selection parameter. Two additional parameters are the mutation rate, which is a biological constant, independent of the individual, describing how likely it is for an allele to change in an individual, and the population size. The way in which population size enters is intuitively clear: when the population is very large, the law of large number is relevant in every period, and the process is close to its deterministic counterpart. When it is not, the process is farther from a deterministic process.

One can now proceed in two different ways: one can approximate the process in phenotype space in continuous time or one can use simulation-based inference to formulate predictions. We use both. This Markov process has an invariant distribution on the polygenic score, i.e., the genetic component of the phenotype. One can study its properties, and check that the difference between the predictions of the model which uses the genotype or the haplotype as variable of interest is not very large.

2.14. Allocation to the two sectors. We now introduce the Roy model component of the model. We may think of a two-dimensional diagram, where the x -axis reports the genetic component of the relevant phenotype, a productive skill. We can then consider two functions of this quantity, corresponding respectively to the productivity of foraging and productivity of farming for each value of this productive skill. For example, these two functions may be, as standard in selection theory, Gaussian functions. These two technologies may coexist at a point in time: in this case, some people

choose farming and others choose foraging, depending at least in part on whether each individual is more productive in one or the other.

The more productive an individual is, the more likely he is to have children. By choosing the activity with the largest productivity one is also increasing the number of descendants. With two different functions describing the productivity depending on the phenotype, there is going to be typically a crossing point of the two productivity curves. Each individual chooses foraging up to a threshold level, and farming for value of the phenotype larger than the threshold level.

I propose the following hypothesis, which I will then test. In the period before 12000 years ago, before the warming of climate, productivity of agriculture was shifting upward, but was overall low because the temperature was too low for a productive agriculture. With the changes following the end of the Younger Dryas, the situation changed. The key observation is that the climate change can make a part of the population more productive because farming becomes more productive, if the climate change increased productivity of farming more than that of foraging.

One way to proceed in order to test the model is to approximate it with a continuous time model. In this case the crucial point is that the speed of change of an allele is determined by the parameter ω times the GWAS coefficient β . In the continuous time approximation the process derived is determined by the product of these two parameters. If you ignore for the moment the variance part, you get an ordinary differential equation. It has a simple exponential solution which gives you a rough idea of how the process evolves. You can then estimate this model as the solution of this ordinary differential equation plus a random term with variance that depends on the frequency.

The approximation is reasonably good. One can then run a simple non-linear least squares model estimation, and check whether the parameter ω is significantly different from zero. The relevant statistic is highly significant and positive compared to the standard error. One can show that the approximation is good, and we can therefore conclude that the ω is significant and positive. Thus one can conclude that there was a selection operating in the direction of alleles that were favorable to farming.

This model has two parameters, that is, the intensity of the selection process and the population size. One can perform the estimation using the simple model, or one can do it using the full model through simulation-based inference. In both cases one gets exactly the same answer.

2.15. A summary. Let us summarize our analysis of the transformation of the distribution of genotypes of the population in the past. We defined a change in a population as a change in the distribution of alleles in that population. We formulated the hypothesis that changes in the distribution of alleles were the outcome of a climate change which made farming more productive than foraging. The allocation of the population between foraging

and farming changed because foraging was less productive than farming, and the larger productivity of farming increased the fitness and thus the number of people who had adopted farming, and thus increased the frequency on alleles that were favorable to agricultural production. Human species underwent a transformation similar to the one that happened as wolves changed into dogs.

It is important to note that the speed of the selection is in scale with observed phenomena. We had 560 generations available to explain the changes, and the parameters of the model are given by biology. We had to explain a change of the phenotype that in terms of normalized quantity goes from a specific value to another. The model gives the variance of the likelihood, and the probability that this happens by chance with no selection is very low.

3. BIOLOGICAL FOUNDATION OF ECONOMIC CHOICE

In this section we will consider the analysis of the choice process, in animal and human subjects, and clarify in specific applications how our approach differs from the traditional one used in economics.

The research program we will develop here is to take the *as if* models of decision making and replace them with mechanistic models. A mechanistic model is defined to be a model that gives a complete description of the process producing data, and that only uses variables in the theory corresponding to quantities that can be measured in the experimental situation. A model of this type is described completely in section (3.5) below.

Our ultimate aim is to replace the hypothesis of the optimality of the choices an individual makes, that is, that the *outcome* of the of choice activity is optimal, with the idea that the *process* that produces choice is optimal. This alternative view considers the process of choice as information acquisition and processing, rather than simply the selection of the best option; and thus considers the natural possibility that there are constraints imposed on this process. Our general claim is that the process, once these constraints are taken into account, is optimal.

Finally, we want to frame our theory of human behavior in economic and strategic situations as part of the general theory of personality.

3.1. The choice process in animal studies. We derive our data from a paper analyzing the choice process (Padoa-Schioppa and Assad (2006)). The paper reports results of an experiment in which Capuchin monkeys are the subjects.

Each monkey in the experiment was trained to follow instructions detailing how the animal could obtain rewards. Since it is impossible to communicate instructions to monkeys, they had to be trained for a long period (approximately six months); the training had to rely on trial and error. In this process, they learned the following rules.

They had to direct their gaze on a fixation point on a screen for one and a half second. Then, two set of squares appeared, one on the left and the other on the right of the fixation point. The two sets of squares were in different numbers and color. The number of squares indicated the quantity of the drink (two or four milliliters, say); the color indicated the type of juice.

For example, four blue squares on the right and one yellow square on the left was a way to tell the monkey “On the right you have four milliliters of peppermint tea, and there on the left you have one milliliter of grape juice”. A monkey typically likes grape juice more than peppermint tea, so when equal quantities of both were offered, the subject would typically choose the grape juice. With the quantities indicated in our example, the choice is more complicated because quantity and quality go in the opposite directions.

After a randomly variable delay (between 1 and 2 s) two dots on the side of the fixation dot appeared next to each of the set of squares. The appearance of these saccade targets are a way to tell the monkey that a choice could be made.

The animal could now fixate one and only one of these dots, for a prescribed length of time. Fixating the dot on the left was a way to tell the experimenter that it had chosen the grape juice. The monkey has to stare at this dot for three quarter of a second, after which it received one milliliter of grape juice. Fixating the square on the right delivered four milliliters of peppermint tea. The task may appear complex; but as we were mentioned the subjects were extensively trained, and they came to the experiment thirsty so they were strongly motivate to pay attention to the process.

You can summarize a trial, like the one we just used as example, with the numbers 1 and 4, because the monkey was offered one unit of grape juice and four units of peppermint tea. We denote this with one of A and four of B , and write $(1A)$ and $(4B)$ for brevity.

There are many different possible menus that we may offer. As we mentioned, the monkey prefers the good A to B if they are offered in equal quantities. But the quantity offered of each good is also important: if B , which is peppermint tea, is offered in sufficiently large quantities, then they will choose the B option.

If one normalizes the utility of the monkey one can see that preferences approximately follows this linear rule: one unit of A is equivalent to four of B . For a different individual monkey the quantity and quality trade-off may be different (three instead of four to be indifferent to one, for example). The exchange value may also depend on special conditions (for example on the animal being more thirsty –which would shift the balance in favor of the B option).

The specific nature of the A and B good changes over blocks. We will however adopt the convention that the A good is the one that, when the quantities are equal, is preferred to B .

3.2. Behavioral and neural data. If you elicit the choices of the monkey over a large number of choice instances you generate data on the frequency of the A versus B choice. As the quantity B is increased and A is decreased one can expect the proportion of the B choice to go up. One can then run a logistic regression on these data and estimate a coefficient representing the strength of the relative preference for A over B .

What we have described so far is the behavioral side of the study: as we study their choices, we learn the relative preference of each individual monkey for one good (e.g. grape juice) versus the other, how much variability there is over trials (for instance, how many times when faced with the choice between 1 of A and 3 of B the monkey chooses the first option), how much time is needed to reach the decision, how this time depends on the difference in estimated utility between the two options, and so on. Of course, one can do the same experiment with humans and elicit the utility function of the human subject.

When subjects are monkeys we can gather data which are considerably much more informative: we can insert an electrode in the brain of the monkey and get a detailed record of the firing rate of the neurons in different regions of the brain. Information processing in the brain is implemented through computations performed by a network of neurons communicating with each other. Communication among neurons occurs through electrical impulses and neurotransmitters.

A neuron fires following approximately a Poisson process. The rate of this Poisson process, called firing rate, depends on the variables that we are studying. In an extremely simplified, but substantially correct, way we may think of the firing rate as a measure of the utility to the individual of the good that the individual is considering. Thus we have information about the firing rate of the neurons that have activity correlated with features we are studying.

These data on firing rates and their dependence on the other circumstances are the information that, in addition to the behavioral data, will allow us to test some interesting hypothesis on how the choice process is organized, and its properties. We will proceed in two steps. The first part of our task is to understand how choices are made given the experiment, subjects and data. The second part is to understand in what sense it is the process of choice, rather than the choice outcome, that is optimal.

3.3. How the choice process is organized. A researcher can get recording of activity from randomly selected neurons in many brain regions, and in particular in the ventro-medial Pre-Frontal Cortex (*vmPFC*) of the brain of the animal being studied. We pay attention here in particular to the *vmPFC* region because large evidence, clinical and experimental, suggests that this is a region likely to be involved in the choice process.

Data mining analysis of the data reveals three clusters, or types, of neurons. By “type of neuron” we mean here a set of neurons for which a specific association between behavioral data and firing rates holds.

The behavioral data in isolation (without the association with neural data) show what one can expect: if we increase the quantity of the good B offered, while keeping constant the quantity of the good A offered, then for large enough quantities of B , option B is chosen. If we consider the association between three types of information (the quantities of A and B offered in a trial, the probability that A or B is chosen, and the firing rate of the neuron under consideration) a pattern emerges.

A first type of neurons appears to code the *value of a specific good offered* (panel (a) of figure (2)). In these neurons, the firing rate increases as the quantity of B offered in the trial increases: it is low when only 1 unit of good B is offered, higher with 2 and even higher with 3. The quantity of A offered on the same trials does not have an effect. In summary, the firing rate for these neurons follows the quantity of B . Other neurons follow instead the quantity of A . We classify this as a type of neuron, recording the quantity of one good being offered, and only on one of the two.

A second type of neuron appears to code the *value of the choice made* (panel (e) of figure (2)): the firing rate goes from high to low as the relative quantity of A offered drops, in those trial in which A is chosen. Analogously, the firing rate goes up again as the quantity of B offered increases, and B is chosen.

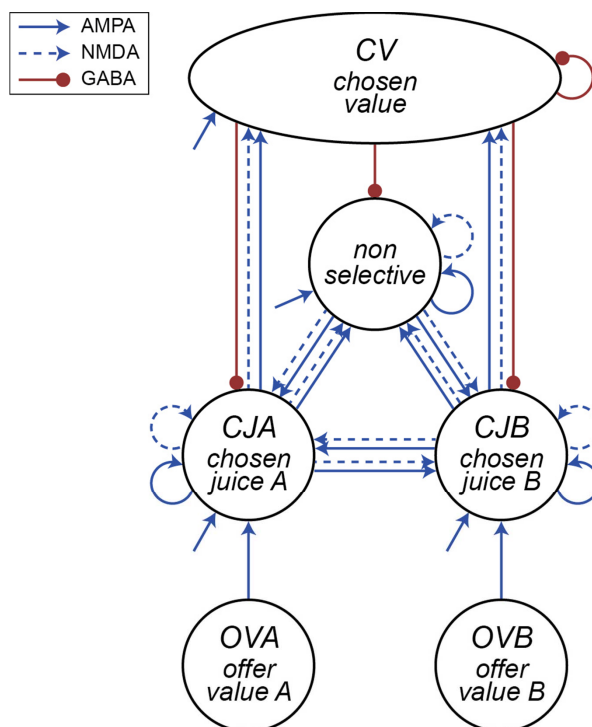
A third and final type of neuron (panel (c) of figure (2)) is associated with the *type of good* (A versus B) which has been chosen in the trial. The firing rate is high and approximately constant when one of the goods is chosen (say A), and low when the other (say B) is. The high firing rate is approximately independent of the value of that good in the trial.

This list is approximately exhaustive: every time you observe the record of the firing rate from a neuron, that neuron will behave in one of these three ways. We want to construct a model giving account of such clustering of neurons in types.

3.4. A template: perceptual decision making. There is a wide consensus that choice in visual perception is organized in a network of neurons that has a specific structure. We can build our analysis on the template offered by models of choice in visual perception, and extend them to economic choices. Let us first briefly consider visual perception tasks.

A task of this type, widely used experimentally, is the following. An animal subject observes random dots on a screen that appear and disappear over time with some regularity. As one watches this process, the dots seem to move randomly, but with higher probability in one of two directions, right or left. The subject, as it observes dots, has to decide whether they are more likely to move to the right or to the left.

FIGURE 1. **A network model of choice.** The cluster of *non-selective neurons* may be ignored. The network presented here is the one relevant for economic choice; as we indicate in the text, with small adjustment the figure can be taken to describe the network relevant for visual perception tasks. In considering the diagram in figure (1), the reader should mentally replace the two goods (indicated in the figure by the letters *A* and *B*) with the two directions (right and left) which are relevant for the visual perception task.

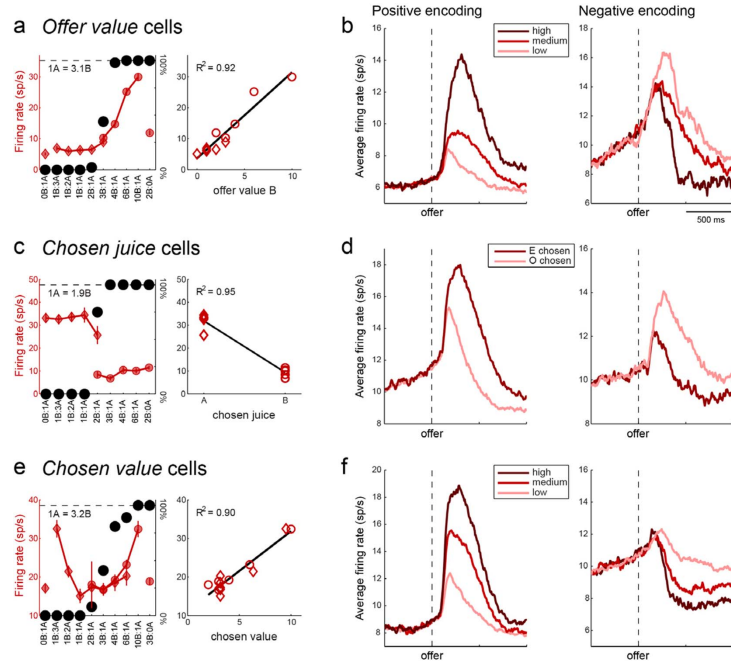


We now proceed to describe the network of neurons that implement choice in visual perception. The network implementing the choice in this task has a very similar structure to the one described for economic choice; so we refer the reader to such network, presented in figure (1).

The two clusters in the bottom tier (indicated in the figure (1) as the *OVA offer value A* — and respectively *B*— cluster) represent a sub-network of neurons coding evidence for left and right respectively. An arrow indicates excitatory impulses — neurons in, say, the left cluster in the bottom tier

excite neurons in the left cluster in the middle tier (indicated in the figure (1) as the *CJA chosen juice A* cluster; respectively *B*). This excitation is proportional to the input that they have received. These neurons in turn activate the downstream neurons in the middle. Then these neurons activate a common pool which are called interneurons (indicated in the figure (1) as the *CV chosen value* cluster). The neurons on the top cluster in return act on the ones in the middle part in an inhibitory function (indicated by the segment terminating with a circle) so the stronger the impulse from the former are, the weaker is the activation of the latter.

FIGURE 2. **Observed data.** Left two panels: pattern of firing of the three types of neurons. Right two panels: firing patterns over time. Top panel: Offer Value cells. Middle panel: Chosen Juice cells. Bottom panel: chosen value cells.



This completes the description of the structure of the network. We could now describe how this network produces choices in visual perception tasks. But we will do this in the case that is of interest to us, that of economic choice.

3.5. A mechanistic network model of economic choice. Let me provide some intuition on the working of the network. Keep in mind that we want ultimately to predict a choice. Let's say that the good A , in the quantity offered in the trial, is worth five units of utility and B is worth only one. Note we are using here quantities of utility, rather than quantity of the goods: we are assuming that the conversion from quantity and type of good into a common currency (utility) has already been made, and the input to the network are provided in these quantities.

Choice is the simple task of selecting the option worth five instead of that worth only one. This is a seemingly elementary task; and in fact economists so far thought that such an elementary task did not need any explanation. But the task has a difficulty: once you consider its neural implementation, this operation of choosing the maximum has to be done using a neural network. We have to describe how the choice of the maximum of two values (five and one in our example) is done by a network.

Let us follow the evolution over time of the activity of neurons in the network, to see how this simple task is executed. It is important to note that each sub-task we describe below can be easily implemented by a network of neurons. For example, adding two values can be easily implemented by a network in which a cluster receives inputs from up-stream clusters, each reporting the value of one of the addends in the sum.

First, neurons in the bottom left cluster are firing at firing rate corresponding to a value of five, while those on the right are firing at a rate corresponding to a value of one. The neurons in each of these clusters in turn are giving impulses corresponding to five and one to the neurons in the middle left and middle right clusters respectively. These clusters give impulses of five and one to the top single cluster. The top cluster now finds the average of these two, which is three, and sends back inhibitory impulses worth three to each of the middle left and right clusters. The middle-left cluster now have incoming impulses worth five, and inhibitory impulses worth three, a net value of two; similarly the middle-right one receives impulses for a net value of negative two. As these effects accumulate over time, the former will converge to maximum activation while the latter will die off. When the activation of the first passes a threshold, the option corresponding to that cluster is chosen. That is how you implement this very simple choice in network of neurons.

We can now go back to the description of the three types of neurons, with patterns of activation and the link with choice as described in section (3.3), their relative location in the network (as presented in figure (1) and their pattern of activation within a trial as hypothesized in this section. There is a perfect match of the three. This is a mechanistic model of choice.

3.6. Gating variables model. The network we analyzed is the core of the process producing choice. We can now examine more precisely the micro-foundation of the way in which the network operates. Presenting this is

necessary to understand an important feature of the choice process, adaptive coding.

Here is a succinct way to present a model of this process. In the model we are using (see Wang (2002), Wong and Wang (2006), Wang (2008)) the activity of a neuron is approximated by a variable describing the fraction of gating variables (NMDA receptors) which are open. The reason why this approximation is legitimate is complex and too technical to be explained here; the reader is referred to Wong and Wang (2006) for the details.

We let S^i be the fraction of gating variables associated to good i which are open. This fraction follows a dynamic over time that has two components. First, a decay proportional to the current activity (the term $-\frac{S^i(t)}{\tau}$), and second a term proportional to the complement of to 1 of $S^i(t)$, namely $(1 - S^i(t))$, multiplied by a term that depends on the rest of the system. Putting all together we get:

$$\frac{dS^i}{dt}(t) = -\frac{S^i(t)}{\tau} + (1 - S^i(t))\gamma\phi(X^i(t)),$$

for $i = A, B$, where for $i = A, B, j \neq i$. The term $X^i(t)$ is defined by:

$$(1) \quad X^i(t) \equiv J^{ii}S^i(t) - J^{ij}S^j(t) + I^i(t)$$

where we have denoted with J^{ij} the coefficients, with the restriction $J^{ij} > 0$, which describe the connection between i and j , and where I^i is the input for the option i . The rate of activation for good i in equation (1) is determined as a sum of a first term, a deterministic quantity declining in proportion to the activation. The second term on the right hand side reduces the excitatory impulse by an inhibitory factor.

Now we have a model and we can predict at the same time on the basis of this model both behavioral data and neural data.

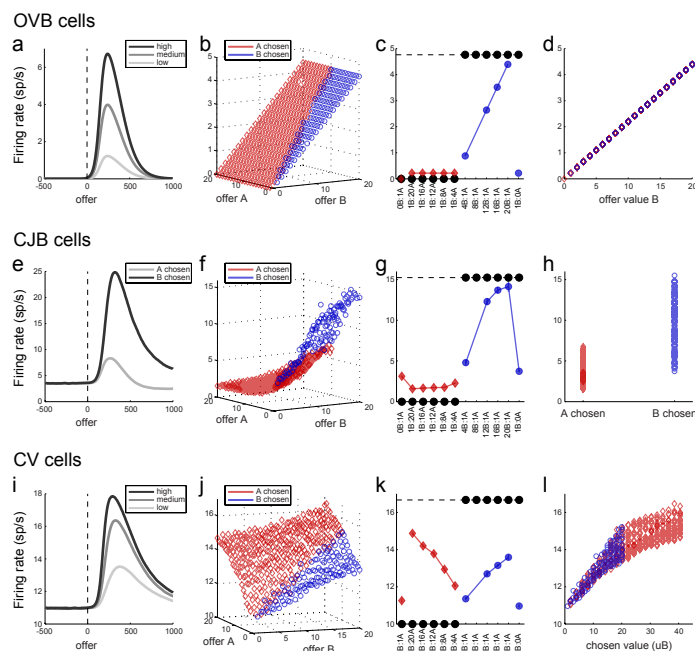
3.7. Choice and firing patterns in experimental data. The predicted neural data correspond very well to the experimental data (Rustichini and Padoa-Schioppa (2015)). Recall that our model is a system of stochastic differential equation. It generates a pattern over time in each of the five clusters of neurons, which we can use to compare to the observed paths, and so we can test how good the match is. Note that we both predictions and observations describe patterns of firing rates and behavioral choices.

The predicted data are reported in figure (3); the goodness of the match can be estimated from the comparison of figures (2) and (3). We are able to match the probability of choice, the pattern of activation of different neurons, and even the pattern over time.

3.8. The micro-foundation of the process. We can now show that the model can explain important characteristics of human economic choice. We want to prove theorems like this:

“We have a cardinal utility because neurons communicate in an additive way”

FIGURE 3. **Predicted data.** Top panel: Offer Value cells. Middle panel: Chosen Juice cells. Bottom panel: chosen value cells.



or like this:

“We have adaptive coding because the spike process of neurons is Poisson.”

We are going first to review very briefly the simple process explaining how a single firing of a neuron is determined. The reader can see the details more extensively in classical textbooks (see for example Dayan and Abbott (2001), Renart et al. (2003)).

The state of a neuron at a point in time is described by its voltage $V(t)$ at time t . The voltage evolves according to the following equation of *leaky integrate-and-fire (LIF)* neuron, in response to stochastic synaptic inputs. The equation gives the new voltage after a small time interval dt :

$$\begin{aligned}
 V(t + dt) - V(t) &= -g(V(t) - V_L)dt \\
 &+ J_E(N_E(t + dt) - N_E(t)) - J_I(N_I(t + dt) - N_I(t))
 \end{aligned}$$

In this equation, V_L is the resting potential of the cell in absence of inputs, and is approximately equal to $-70mV$. The parameter J measures the strength of the synapse (more precisely, J is the amount of positive charge

entering the membrane due to the spike). The N 's denote the Poisson processes of the number of spikes of the input neurons. The subscripts E and I refer respectively to excitatory or inhibitory inputs.

The variable $V(t)$ follows this equation until threshold V_{th} is reached. When this happens, the neuron fires, and the voltage is reset to a V_r value. The process then starts again, with this new initial condition and dynamics described by the above equation. This dynamical system reduces the choice process to its elementary components. We want to explain adaptive coding with just these elements, and thus prove:

Theorem We have adaptive coding **because** the spike process is Poisson

Before we proceed, we need to summarize what adaptive coding is, and what we know about it.

3.9. Adaptive coding. Adaptive coding describes the property of a neural system coding rewards of adapting to properties of the environment. The term adaptive is used because the responses of the system adapt to those of the environment in which the individual operates. Clear demonstrations of adaptive coding are given in (Tobler et al. (2005), Padoa-Schioppa (2009), Diederer et al. (2017); the theory problem posed by adaptive coding was considered first in Padoa-Schioppa and Rustichini (2014)). Adaptive coding is related to the phenomenon of reference dependence studied in economics: the two terms describe ultimately similar phenomena, the dependence of choice from reference values in the environment.

To illustrate the neural side of adaptive coding we refer to Padoa-Schioppa (2009), in particular the figure 5 of that paper, which describes the neural response measured by the firing rate as function of the quantity offered of a good. Obviously, as the quantity increases the firing rate increases, in approximately a linear fashion. What is crucial however is that the slope of this linear response depends on the range of the offer.

The figure 5 cited describes the firing rate of the five types of neurons that we described earlier. All of them exhibit adaptive coding. The steepest lines correspond to an experimental environment in which you are offered between a minimum of zero and maximum of two quantities of a good. Given different quantities offered, you can trace the firing rate of a neuron. Other lines, considerably flatter, correspond to a different experimental environment in which you are offered between the minimum of zero and the maximum of approximately 9 to 10.

Averages taken for each experimental condition (shown in the top right panel of figure 5 of Padoa-Schioppa (2009)) clearly show that the adaptation of the slope is such that the firing rate is always normalized in such a way that, in spite of the fact that the range of rewards is different in different experiments, the range of the firing rate is always the same. The maximum is always nine to ten spikes and the minimum is always zero. Note that the baseline firing rate has been subtracted, so what we are considering

here is actually the difference from the minimum firing rate, which is then automatically normalized to zero.

These data show the existence of adaptive coding, that is of the fact that the firing rate of a neuron which is coding rewards (or is coding features in visual perception) adjusts so that the dynamic range, which is the difference between the minimum and the maximum firing rate, is approximately the same in all environments (nine to ten spikes in our case). If the quantities of good offered all fall in the interval between zero and two, you observe a firing rate of zero at zero and that of nine at two; if you offer between 0 and 10 you also observe a firing rate of zero at zero and that of 9 at 10. Why? We examine this question in (Rustichini et al. (2017)).

3.10. Why adaptive coding? Let us consider again the visual perception case, which is familiar in everyday experience. As we already noted, when you move from one room to another with different, say larger amount of light, your perceptual system adjusts, by making the response proportionally weaker; after this initial period of adjustment, the response stabilizes. We are investigating the following question: when the monkeys evaluate rewards, do they behave in the same way as we do when you are put in a bright or in a dim room? The answer seems to be positive. The feature common to reward coding and to visual perception is that the slope of the firing rate changes so that in each environment the rate varies from its minimum value to the maximum, thus exploiting the full dynamic range. That's the regularity we observe in this data.

The question is why this adjustment occurs. What we need, to answer this question, is a key property of the process we are studying that makes this choice of re-scaling, with respect to the distribution of rewards in the environment, optimal.

Let's recall that the network described earlier is stochastic (because it is a combination of Poisson processes) and it has this adaptive coding property. We now claim that the re-scaling of adaptive coding is optimal; and we are going to show why this is so. We will see that this perhaps surprising property only holds when coding is performed by Poisson processes.

A neuron always fires at least five to six hertz, that is five to six number of spikes per second; the average firing rate in the ventromedial prefrontal cortex is about 16 spikes. How does a Poisson process look at a rate of 16? Simple computations show that it is approximately like a normal distribution with a mean of 15, which is also the variance by the well-known property of the Poisson process.

3.11. Implications of these properties. In summary, we have a process where you can adjust the mean firing rate, but when you adjust the mean, you are also adjusting the variance. Let us now consider a very simple problem, to understand how the fact that neural coding follows a Poisson process explains in turn the existence of adaptive coding.

I offer you to choose from a menu of two options, delivering respectively utility x and y , quantities that are unknown to you. You have to choose one of the two options, and you will receive the corresponding value, x or y , so you want to choose the maximum of the two quantities. You do not observe them directly, but you can observe a noisy signal associated with them, as follows.

First, there is a random variable X , associated with x , which is normally distributed with mean sx and variance also sx , where s is a real number, which corresponds in the complete model to the slope of the firing rate. The parameter s is under your control. For y there is a variable Y which is constructed similarly.

You will not observe the individual realization of these quantities, but the difference of the two, $X - Y$ say. By the symmetry of the problem, once you receive the signal, you will choose option giving the unknown quantity x if the signal is positive, and y otherwise. You can affect the signal before you receive it, and your choice variable is s , the slope of the firing rate.

The simple example is an appropriate simplification of our original problem, because a the Poisson process is approximately normal, with an approximation reasonably good in the range which is relevant for our purposes, namely when the number of spikes is about 10 or above. The really crucial property of Poisson processes that we are using here is that the mean and the variance must be the same.

A simple algebraic rearrangement reduces the problem of maximizing the probability of making the high value choice (choosing x when $x \geq y$) to that of maximizing the probability that a standard normal is larger than $-(x - y)\sqrt{s}$. Therefore, if you can choose s you should choose it as large as possible within the feasibility constraints, for example less than a maximum physiologically possible.

We have shown that the real reason why we have reference dependent choices is that we need to adjust our choice process to the properties of the environment to maximize expected payoff, and this property follows from the Poisson structure of the neural coding.

4. CONCLUSIONS

The ideas presented in these three lectures belong to a long standing tradition in philosophical thought. It is useful to discuss here this tradition, and possible alternative philosophical positions, to place them in the proper perspective.

In the lectures we developed the thesis that social, economic and political behavior should find an explanation in deeply rooted, ultimately biological, characteristics of individuals, and that properties of the societies in which these individuals live should be inferred from the individuals' characteristics. This view, which we may call materialistic, is opposite to an idealistic thesis

that properties of society are in the final analysis determinants of individual characteristics.

This fundamental premise goes back at least to Thomas Hobbes, and before him to Niccolò Machiavelli. The *Leviathan* has political philosophy as the field of examination, but it begins with an extended analysis of the fundamental motivations of an individual *in isolation*. The treatise develops a method which deduces the properties of the political equilibrium from the elementary properties of human nature, and so it begins with a detailed examination of these elementary properties. It proceeds from an analysis of characteristics of an individual in isolation, establishes their main features, and then goes on to show the implications of these facts for the interaction of these individuals in social life. The analysis progresses from Sense, to Imagination, Speech, Passions (chapters 1 to 9) going then on to the socially relevant emotions: Power, Worth, Dignity. The analysis culminates in chapter 13, where Hobbes takes as the premise the observation that the difference among human is small to conclude what the main properties of a social organization will be. “Small” here has a very precise meaning, because the distance scale is well defined: any individual has enough cunning or physical strength to threaten the others. Thus, the analysis of the premises has led to the conclusion that, lacking an agreement among humans, conflict is the natural state in society. Ultimately, the fundamental human motives are desire of power, made possible in everyone by the fundamental equality of the threat to others, and fear of death.

The idea that an analysis of human events, and of good political arrangements, must begin from the understanding of human nature, rather than from what desirable properties these arrangements should have, is the fundamental insight of Machiavelli. Leo Strauss, contrasting the classical and modern school of political philosophy indicates Hobbes, and Machiavelli before him as founders of the modern school.

Let us recall the contrast, set up in the works of Strauss (in particular Strauss (1957)), between the classical and modern thought. The classical school in political philosophy (first among others, Plato and Aristotle) had a clear purpose: “Classical political philosophy is guided by the question of the best regime.” This statement requires a definition of best regime, and this is also unambiguous for Strauss: “the aim of human life, and hence of social life, is not freedom but virtue.” On these grounds, “the classics rejected democracy” because “Freedom as a goal is ambiguous, because it is freedom for evil as well as good.”

The modern school stands in contrast, and “the founder of modern political philosophy is Machiavelli. (Strauss (1957)).” Echoing the famous passage of *Il principe*¹ Strauss identifies this fundamental difference between classical and modern thought:

The traditional approach was based on the assumption that morality is something substantial: that is is a force in the soul of man, however ineffective it may be especially in the affairs of states and kingdoms. Against this assumption Machiavelli argues as follows: virtue can be practiced only within society.

The statement “morality is a force in the soul of man” is worth noting well. This is a factual statement, not a difference in value, even less a difference in methodology. Thus, the difference between classical school and modern school relies on a difference in the substantial view of human nature and its very existence.

The view of economics presented in these lectures is the expression of one of the two sides in the contrast between classical and modern thought. We look at human events firmly from the modern side. In contrast, economics as a science seems to be moving slowly but decisively in the direction of making abstract virtues, rather than reality of human nature, the guiding principle of its analysis and proposals.

This process moves in two main directions. The first direction is to deny the ability of individuals, in general, to be able to act effectively in their own interest. The second direction is to identify moral principles as innate and frequently contradictory to self interest. But I believe that if indeed altruism, fairness or aversion to inequality are “forces in the soul of men” then these forces will find a way to in social life, as expression of the free choice of individuals. This outcome however requires making freedom, rather than virtue, the guiding ethical principle.

The outcome of this transformation of economics would not be good, because it would deprive economics of its main strength when compared to the other social sciences, which is the realism of the analysis. These lectures are trying to help to reverse this course.

¹..it being my intention to write a thing which shall be useful [...], it appears to me more appropriate to follow up the real truth of the matter (*andare dritto alla verità effettuale della cosa*) than the imagination of it; for many have pictured republics and principalities which in fact have never been known or seen, [*Il Principe*, Chapter 15]

REFERENCES

- BAULT, N., M. JOFFILY, A. RUSTICHINI, AND G. CORICELLI (2011): “Medial prefrontal cortex and striatum mediate the influence of social comparison on the decision process,” *Proceedings of the national Academy of sciences*, 108, 16044–16049.
- BECKER, G. AND N. TOMES (1979): “An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility,” *Journal of Political Economy*, 87, 1153–89.
- BERG, J. J. AND G. COOP (2014): “A population genetic signal of polygenic adaptation,” *PLoS genetics*, 10, e1004412.
- BOCQUET-APPEL, J.-P. (2011): “When the world’s population took off: the springboard of the Neolithic Demographic Transition,” *Science*, 333, 560–561.
- BOCQUET-APPEL, J.-P. AND O. BAR-YOSEF (2008): *The Neolithic demographic transition and its consequences*, Springer.
- BOCQUET-APPEL, J.-P. AND A. DEGIOANNI (2013): “Neanderthal demographic estimates,” *Current Anthropology*, 54, S202–S213.
- BOCQUET-APPEL, J.-P. AND C. MASSET (1982): “Farewell to paleodemography,” *Journal of Human Evolution*, 11, 321–333.
- BOWLES, S. (2011): “Cultivation of cereals by the first farmers was not more productive than foraging,” *Proceedings of the National Academy of Sciences*, 108, 4760–4765.
- BOWLES, S. AND J.-K. CHOI (2013): “Coevolution of farming and private property during the early Holocene,” *Proceedings of the National Academy of Sciences*, 110, 8830–8835.
- COMMITTEE ON HIGHER EDUCATION; GREAT BRITAIN (1963): *Higher education: report of the Committee appointed by the Prime Minister under the chairmanship of Lord Robbins 1961-63 [Robbins report]*, Her Majesty’s Stationery Office, London, England.
- CORICELLI, G. AND A. RUSTICHINI (2010): “Counterfactual thinking and emotions: regret and envy learning,” *Philosophical Transactions of the Royal society B: Biological sciences*, 365, 241–247.
- DAWES, C. T., A. OKBAY, S. OSKARSSON, AND A. RUSTICHINI (2021): “A polygenic score for educational attainment partially predicts voter turnout,” *Proceedings of the National Academy of Sciences*, 118, e2022715118.
- DAYAN, P. AND L. F. ABBOTT (2001): *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, MIT Press.
- DIEDEREN, K. M., H. ZIAUDEEN, M. D. VESTERGAARD, T. SPENCER, W. SCHULTZ, AND P. C. FLETCHER (2017): “Dopamine modulates adaptive prediction error coding in the human midbrain and striatum,” *Journal of Neuroscience*, 37, 1708–1720.
- DISNEY, E. R., I. J. ELKINS, M. MCGUE, AND W. G. IACONO (1999): “Effects of ADHD, conduct disorder, and gender on substance use and

- abuse in adolescence,” *American Journal of Psychiatry*, 156, 1515–1521.
- FEDDERSEN, T. J. AND W. PESENDORFER (1996): “The swing voter’s curse,” *The American economic review*, 408–424.
- FRIEDMAN, M. (1953): *The methodology of positive economics*, Chicago.
- IACONO, WILLIAM ND CARLSON, S. R., J. TAYLOR, I. J. ELKINS, AND M. MCGUE (1999): “Behavioral disinhibition and the development of substance use disorders: Findings from the Minnesota Twin Family Study.” *Development and Psychopathology.*, 11, 869–900.
- ICHINO, A., A. RUSTICHINI, AND G. ZANELLA (2022): “College education, intelligence, and disadvantage: policy lessons from the UK in 1960-2004,” Tech. rep., CEPR Discussion Paper 17284.
- JANDAROVA, N. AND A. RUSTICHINI (2023): “Selection and the Roy Model in the Neolithic Transition,” *Available at SSRN 4357095*.
- KINGMAN, J. F. (1982): “On the genealogy of large populations,” *Journal of applied probability*, 19, 27–43.
- MARX, K. (1875): *Critique of the Gotha program*, PM Press.
- MARX, K. AND F. ENGELS (1845): “The German ideology: Critique of modern German philosophy according to its representatives Feuerbach, B. Bauer and Stirner, and of German socialism according to its various prophets,” *Marx, K. & Engels, F.(1976a). Karl Marx and Frederick Engels Collected Works*, 5, 1845–47.
- MCMURRAY, J. C. (2013): “Aggregating information by voting: The wisdom of the experts versus the wisdom of the masses,” *Review of Economic Studies*, 80, 277–312.
- PADOA-SCHIOPPA, C. (2009): “Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex,” *Journal of Neuroscience*, 29, 14004–14014.
- PADOA-SCHIOPPA, C. AND J. ASSAD (2006): “Neurons in the orbitofrontal cortex encode economic value.” *Nature*, 441, 223–226.
- PADOA-SCHIOPPA, C. AND A. RUSTICHINI (2014): “Rational Attention and Adaptive Coding: A Puzzle and a Solution,” *American Economic Review*, 104, 507–13.
- PROTO, E., A. RUSTICHINI, AND A. SOFIANOS (2019): “Intelligence, personality, and gains from cooperation in repeated interactions,” *Journal of Political Economy*, 127, 1351–1390.
- (2022): “Intelligence, Errors, and Cooperation in Repeated Interactions,” *The Review of Economic Studies*, 89, 2723–2767.
- RENART, A., N. BRUNEL, AND X.-J. WANG (2003): *Computational Neuroscience: A Comprehensive Approach*, Chapman & Hall/CRC, chap. 15, 431–490.
- ROBINSON, S. A., S. BLACK, B. W. SELLWOOD, AND P. J. VALDES (2006): “A Review of Palaeoclimates and Palaeoenvironments in the Levant and Eastern Mediterranean from 25,000 to 5000 Years BP: Setting the Environmental Background for the Evolution of Human Civilisation,” *Quaternary Science Reviews*, 25, 1517–1541.

- ROBSON, A. J. (2010): "A bioeconomic view of the Neolithic transition to agriculture," *Canadian Journal of Economics/Revue canadienne d'économique*, 43, 280–300.
- ROUSSEAU, J.-J. (1999): *Discourse on the Origin of Inequality*, Oxford University Press, USA.
- ROWTHORN, R. (2011): "A bioeconomic view of the transition to agriculture: a comment," *Canadian Journal of Economics/Revue canadienne d'économique*, 44, 1044–1047.
- ROWTHORN, R. AND P. SEABRIGHT (2010): "Property rights, warfare and the neolithic transition," Working paper.
- ROY, A. D. (1951): "Some thoughts on the distribution of earnings," *Oxford economic papers*, 3, 135–146.
- RUSTICHINI, A., K. E. CONEN, X. CAI, AND C. PADOA-SCHIOPPA (2017): "Optimal coding and neuronal adaptation in economic decisions," *Nature Communications*, 8.
- RUSTICHINI, A., W. IACONO, J. LEE, AND M. MCGUE (2023): "Educational Attainment and Intergenerational Mobility: A Polygenic Score Analysis," *Journal of Political Economy*.
- RUSTICHINI, A. AND C. PADOA-SCHIOPPA (2015): "A neuro-computational model of economic decisions," *Journal of Neurophysiology*, 114, 1382–1398.
- SCHULTZ, W., P. DAYAN, AND P. R. MONTAGUE (1997): "A neural substrate of prediction and reward," *Science*, 275, 1593–1599.
- STERN, A. J., L. SPEIDEL, N. A. ZAITLEN, AND R. NIELSEN (2021): "Disentangling selection on genetically correlated polygenic traits via whole-genome genealogies," *The American Journal of Human Genetics*, 108, 219–239.
- STRAUSS, L. (1957): "What is political philosophy?" *The Journal of Politics*, 19, 343–368.
- TOBLER, P. N., C. D. FIORILLO, AND W. SCHULTZ (2005): "Adaptive coding of reward value by dopamine neurons," *Science*, 307, 1642–1645.
- WANG, X.-J. (2002): "Probabilistic Decision Making by Slow Reverberation in Cortical Circuits," *Neuron*, 36, 955–968.
- (2008): "Decision making in recurrent neuronal circuits," *Neuron*, 60, 215–234.
- WATKINS, C. J. AND P. DAYAN (1992): "Q-learning," *Machine learning*, 8, 279–292.
- WONG, K.-F. AND X.-J. WANG (2006): "A recurrent network mechanism of time integration in perceptual decisions," *Journal of Neuroscience*, 26, 1314–1328.