The Evolution of Personal Standards into Social Norms

Hannes Rusch[a,b], Alexander Vostroknutov[b]

[a] *Max Planck Institute for the Study of Crime, Security and Law, Günterstalstraße 73, Freiburg im Breisgau, 79100, Germany*
[b] *Department of Microeconomics and Public Economics, School of Business and Economics, Maastricht University, P.O. Box 616, Maastricht, 6200 MD, The Netherlands*

Abstract

Social norms have become a conceptual cornerstone in the study of human decision making across the social sciences. The functions of social norms in guiding individual and collective decision-making have been extensively scrutinized empirically, too. However, possible evolutionary origins of the psychological mechanisms required to carry out these functions are less well understood. In particular, trajectories from individually adaptive to socially functional heuristics for norm formation have rarely been studied. Here, we trace such a trajectory. We present a model that allows for the comparison of two heuristics broadly applicable across individual and social decision contexts: 'rejoicing' own achievements vs. 'regretting' missed opportunities. We find that (i) both perform better than the homo œconomicus in individual decision problems under plausible ecological assumptions and (ii) each is adaptive in different cost environments. We argue that observation (i) provides a potent microfoundation for social norms as a product of co-optation of individually evolved heuristics, i.e., a reduction of social norm formation to the evolution of individual traits. Moreover, observation (ii) lends itself to empirical testing, thus laying the ground for a new wave of studies in the literature fascinated with human norm psychology.

Keywords: norms, personal standards, evolution, regret, rejoice

1. Introduction

The human eye is an astounding organ. Having evolved long before humans parted ways with the other primates, eyes were certainly indispensable for individual survival under the conditions our early ancestors lived in. Their usefulness in a plethora of behavioral domains is so obvious that no one is really surprised *that* eyesight evolved, the surprises lurk more in *how* it evolved; especially when one realizes that eyesight evolved in countless different ways across the living things [1]. In humans, moreover, the eye not only is extremely useful for the individual, it also serves several social functions including threatening, flirting, pointing directions, or signaling sadness [2]. Thus, the eye is a good example of an individually highly useful product of evolution which was 'co-opted' [3] for social uses later on.

There is no doubt that groups of humans tremendously benefit from being able to coordinate collective action via the use of social norms, i.e., individually available ideas of what one ought to do in a certain situation [4]. An important question then is, of course, how this ability came about. Some philosophers would actually argue that, like eye-sight evolved 'because' there is sunlight that can be detected and used to navigate the physical world, our capabilities for normative thinking evolved 'because' there is a sphere of 'normative truth' or 'objective morals' which humans are able to tap into to navigate the social world [5].

Luckily, ontologically less demanding explanations exist as well. While they also point to the usefulness of conventions and social norms for groups of humans, they do not require the existence of norms independent of human—and possibly also non-human—minds [6], i.e., brains [4, 7, 8]. Still, many of these approaches could be characterized as what we would call 'giant leap' approaches: they tackle the evolution of the psychology required to individually digest social norms under the assumption that those social norms are already somehow 'out there' to be harvested [9]. This creates intricate versions of 'chicken vs. egg' dilemmas, of course.

Here, we take one analytical step back and ask: Are there possible ways for a norm psychology to evolve when there are no social norms 'out there' yet? Our answer is: Yes, there are at least two such ways, a 'regret' and a 'rejoice' heuristic. Our main contribution is that we show how agents equipped with either of these two heuristics for setting *personal standards* of behavior in individual decision problems can fare evolutionarily better than agents without such psychological machinery under only very mild assumptions about the ecology they live in. We then suggest how this machinery, evolved for purely individualistic purposes, could have been co-opted for the construction of personal and social *norms*.

Apart from offering an ontologically undemanding, plausible, and testable explanation for the evolution of human norm psychology, the two heuristics we study are highly generic. They can be applied to individual decision-making problems, which will be our focus in this paper. Beyond this, however, they can easily be applied to social decision-making problems, i.e., games, too; we outline a few applications in Section 6. This qualifies these heuristics as candidates for successful maxims *sensu* [10] thus placing them in an emerging branch of the literature in evolutionary game theory [11].

The related literature that takes an evolutionary approach to the study of the emergence of norms is very sparse—it is comprehensively reviewed in [12]. Importantly, relative to the 'radically individualistic' approach that we take here, this literature is still 'social' in its assumptions about the driving forces behind norm evolution. For example, Calabuig et al. [13] present a quite specific model in which personal standards, $\hat{e}$, are operationalized as individually targeted, non-verifiable effort levels in a team production game. These $\hat{e}$'s then evolve, i.e., gradually change over time, driven by two forces: disutility from deviations between personal standard and actual effort chosen, labeled 'consistency' and being a purely individualistic channel, but also disutility from deviations from the population average effort, labeled 'conformity' and being a social channel. Gavrilets [12] presents a richer and somewhat less specific model whose agents also possess personal norms; however, change of these norms is again partially induced via social channels. In addition to consistency and conformity with peers, Gavrilets includes conformity with some exogenously imposed authoritative norm.

In the present paper, we complement this small but growing literature in at least three key respects: (i) we take a radically individualistic stance in which norm psychology requires no information about peer behavior whatsoever; (ii) by taking a maxim-based approach we impose only very mild structural restrictions on the decision problems faced by individuals; (iii) we demonstrate how heuristics for forming personal standards that are individually adaptive can lend themselves to co-optation for fulfilling social functions, thus indicating a possible solution of the 'chicken vs. egg' dilemmas faced by earlier work on this question.

The remainder of this paper is organized as follows. In Sections 2 and 3 we introduce two heuristics for generating personal standards of behavior:

‘rejoicing’ own achievements and ‘regretting’ missed opportunities. In Section 4 we present an evolutionary model illustrating the superiority of these two heuristics relative to the homo œconomicus benchmark in a specific ‘ecological’ setting of an individual effort task. In Section 5 we prove that this superiority holds quite generally. In Section 6 we discuss how our findings generalize to social settings and conclude.

2. Simple ‘rejoicing’ and ‘regretting’

Literatures in economics and neuroscience have been studying regret and disappointment, defined as feeling disutility from not being able to achieve better feasible outcomes, for quite some time [14, 15, 16]. These studies started from observed behavioral deviations from expected utility theory in decisions under risk (e.g., Allais paradox). They proposed that taking into account better outcomes that could have realized can explain some of these deviations. Later, behavioral and neuroscience experiments provided evidence of the influence of regret and disappointment in risky choices and their manifestation in the brain [17]. It was suggested that these psychological heuristics evolved because they are helpful in learning what could have been achieved and can thus stimulate the decision-maker to strive for more, eventually increasing her expected payoff [18].

Despite this idea, that regret and similar psychological heuristics can enhance survival chances, not being new, we know of no formal theory describing exactly how this happens. The reason for this may be that regret was originally used to explain choices between two lotteries [14]: in that setting it is indeed unclear how regret might be welfare-enhancing in general. Therefore, in this section, we look at the effects of very simple ‘regretting’ and ‘rejoicing’ heuristics in a different setting, namely that of generic individual decision tasks. In these, an agent chooses a level of costly effort and obtains utility in return. This is a standard problem used ubiquitously in many fields of economics to model the relationship between personal costs and exerted effort.

Suppose that the agent chooses how much effort $x \in [0, L]$ to exert in an individual task where the cost of effort is $c(x)$ and the resulting utility is $u(x)$. Then, the standard homo œconomicus agent solves

$$\max_{x\in[0,L]} u(x) - c(x),$$

which, under standard assumptions on $c(x)$ and $u(x)$, generates the interior solution x^* satisfying $u'(x^*) = c'(x^*)$. Now, suppose in addition that the agent's chances of survival are proportional to their effort—e.g., because higher x^* produces more 'wealth' for the agent which accumulates and makes the agent more 'resilient' (think of the agent's nutrition for example). The question we want to ask is whether a modified type of agent who anticipates feelings like regret or rejoice after obtaining $u(x)$ would choose higher effort. In other words: can regret or rejoice push the agent to work more and thus to increase their chances of survival relative to homo œconomicus?

To answer this question, we need to understand exactly how regret or rejoice are computed by the agent. Specifically: whether the agent feels regret/rejoice about the whole combined utility, $U(x) = u(x) - c(x)$, or only about the resulting consumption utility, $u(x)$. We believe that this is not just a matter of making some assumption, but rather the matter of the ecological situation in question. Sometimes, the cost of effort and the consumption utility obtained from it are perceived *simultaneously*. For example, when the agent sits on a tree and eats the fruit that they can reach with their hands. The cost of reaching the fruit and the consumption utility of the fruit are experienced together and it might be reasonable to assume that in this case the agent feels regret or rejoice about $U(x) = u(x) - c(x)$, simply because this combined utility is what the agent perceives in that moment. In other situations, however, the agent might exert effort x first, feel the cost $c(x)$ of it, but obtain the consumption utility not immediately but *with a delay*. Imagine, for example, the agent waiting for prey in an ambush during a hunt: they decide how much time to wait and experiences the cost of waiting before catching the prey, bringing it home, cooking, and then eating it. In such situations, it is plausible that the agent feels rejoice (regret) only about the resulting consumption utility $u(x)$ of the prey they (could have) caught, since the cost of effort is not felt anymore at the time of consumption; at that point, that cost was already 'paid' and the agent's metabolism might have had enough time to recuperate, i.e., the agent already feels 'normal' again when enjoying $u(x)$.

Thus, let us see now what happens when we enhance our agent's utility with simple additional regret or rejoice terms. When $U(x)$ is used for the computation of regret or rejoice, we are in a situation like the one shown in the left panel of Figure 1. Here, for any choice of x, the agent feels rejoice of the size $U(x) - U_0$ (shown as the blue line) or regret of the size $U_L - U(x)$ (green line). These are *simple* regret and rejoice terms, because the agent just compares what they get at x with what they could have gotten in the best situation (U_L) for regret or the worst situation (U_0) for rejoice. Accordingly, the agent's utility function can then be written as $U(x) - \sigma(U_L - U(x))$ for regret and as $U(x) + \sigma(U(x) - U_0)$ for rejoice. Here, we let $\sigma > 0$ be an individual parameter that determines how important simple regret or rejoice are to the agent.

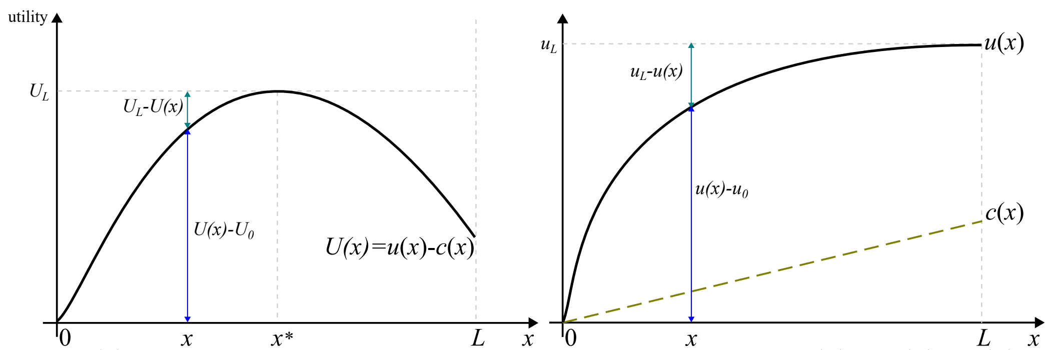

Figure (1) *Left.* Simple regret and rejoice when the whole utility $U(x) = u(x) - c(x)$ is used. *Right.* Simple regret and rejoice when only consumption utility $u(x)$ is used.

An important observation here is that the resulting utility functions $(1+\sigma)U(x) - \sigma U_L$ and $(1+\sigma)U(x) + \sigma U_0$ have maxima at exactly the same x^*, the one that satisfies $U'(x^*) = 0$. In other words, agents with regret or rejoice are choosing the same level of effort as the standard homo œconomicus agent. Thus, in situations where the cost of effort is felt together with the consumption utility we should not expect any differences in behavior between standard agents and agents with regret or rejoice. This also means, of course, that in such conditions agents with simple regret or rejoice do not lose in evolutionary competition with homo œconomicus as long as effort is the only thing that counts.

However, things are very different when costs and consumption are separated in time. Here, as shown on the right panel of Figure 1, the regret and rejoice terms become $u_L - u(x)$ and $u(x) - u_0$ and the utilities of agents with regret or rejoice become $u(x) - \sigma(u_L - u(x)) - c(x)$ or $u(x) + \sigma(u(x) - u_0) - c(x)$, respectively. These can be rewritten as $(1+\sigma)u(x) - c(x) - \sigma u_L$ and $(1+\sigma)u(x) - c(x) - \sigma u_0$. Since the derivative of consumption utility with respect to effort is now multiplied in both cases with the factor $1 + \sigma > 1$, we should expect that agents with simple regret or rejoice will put more effort into the task and thus choose optimal effort higher than x^*, the level chosen by the homo œconomicus. Since we relate the level of optimal effort to survival chances, we should conclude that agents with simple regret or rejoice should outperform homo œconomicus in situations where costs and consumption are separated in time. Thus, we should expect agents with simple regret and/or rejoice to successfully invade populations of homo œconomicus agents and replace them whenever the ecological assumptions we made are satisfied, i.e., whenever there are at least a few situations in which consumption is temporally separated from effort. Moreover, agents who feel more regret or more rejoice, i.e., those with higher σ, also work more than agents with lower σ. So, we should also expect that the strength of feelings of simple regret and rejoice, measured by σ, should gradually increase as agents who feel them more intensely will put more effort and outperform agents with lower σ.

This very simple argument can explain why regret and rejoice could have evolved for purely individual purposes: agents with utility functions that include regret or rejoice work more in individual tasks than the standard homo œconomicus, which gives them a survival advantage. So, in a population of agents indexed by i, we can say that $x^*(\sigma_i)$—the different optimal effort levels induced by individual strengths of feelings of regret or rejoice—represent agents' *personal (working) standards.*

3. Sophisticated 'rejoicing' and 'regretting'

The previous section showed that agents with simple regret or rejoice perform better than homo œconomicus in some individual tasks and thus could win an evolutionary competition with them. However, these simple forms of regret and rejoice do not fully exhaust the possibilities for regretting and rejoicing that the agent has; it could well be that more sophisticated forms of regret or rejoice fare even better, of course.

Indeed, Kimbrough and Vostroknutov [19] suggest that regretting or rejoicing might not be simple, but *sophisticated.* For example, the agent could compute total regret/rejoice as a sum of regrets/rejoices for different counterfactual outcomes. The graphs on Figure 2 illustrate these potential regrets/rejoices with red arrows. In the case of separate utility and cost and K outcomes, we could have an individual regret utility function of the form $u(x) - \phi(\sum_k u(x_k) - Ku(x))$ or individual rejoice utility function $u(x) + \phi(Ku(x) - \sum_k u(x_k))$, where $\phi > 0$ represents the strength of *sophisticated* regret or rejoice. Taking sophistication to extremes produces the utility functions that take into account all possible regrets or all possible rejoices, as illustrated by the green and blue areas on both panels of Figure

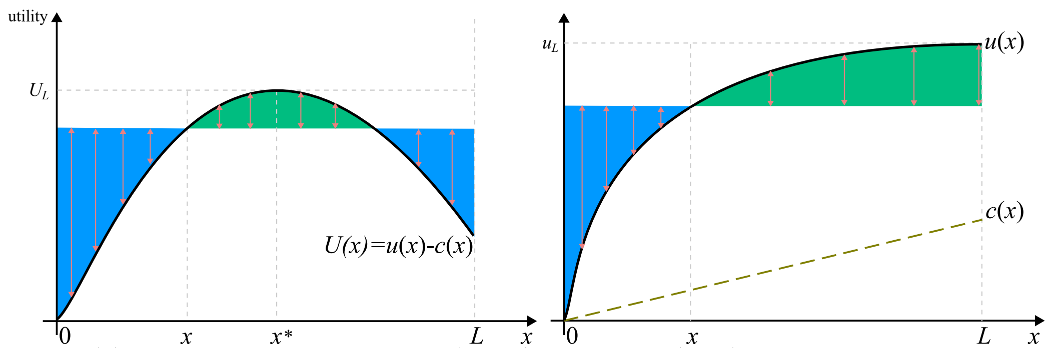

Figure (2) *Left.* Sophisticated regret (green) and rejoice (blue) in the case with combined consumption and cost. *Right.* Sophisticated regret (green) and rejoice (blue) in the case with separated consumption and cost.

2. For the case of separated utility and cost (the right panel of Figure 2), sophisticated regret and rejoice are represented by the two utility functions

$$u_{REG}(x) = u(x) - \phi \int_x^L u(t) - u(x)dt$$

and

$$u_{REJ}(x) = u(x) + \phi \int_0^x u(x) - u(t)dt.$$

To see how well these utilities do in comparison to homo œconomicus and agents with simple regret and rejoice, let us first focus on the case of combined consumption and cost as we did in the previous section. From the left panel of Figure 2 it is clear that agents with sophisticated regret and rejoice will maximize at the same effort x^* as homo œconomicus and agents with simple regret and rejoice. This is so simply because the aggregated regret represented by the green area is the smallest at x^* while the aggregated rejoice (blue area) is the highest at x^*. Thus, agents with sophisticated regret or rejoice do not do worse than homo œconomicus or agents with simple regret and rejoice in individual tasks with combined consumption and costs.

To consider sophisticated regret/rejoice in a more interesting case with separated consumption and cost (the right panel of Figure 2) we use the Fundamental Theorem of Calculus and compute the derivatives of u_{REG} and u_{REJ} given by

$$u'_{REG}(x) = u'(x)[1 + \phi(L - x)]$$

and

$$u'_{REJ}(x) = u'(x)[1 + \phi x].$$

Notice that the derivatives are higher than those for simple regret and rejoice for any $x \in [0, L-1]$ for regret and for any $x \in [1, L]$ for rejoice. Thus, as long as L is large, sophisticated regret and rejoice will induce more work from the agent than simple regret and rejoice. Moreover, the more sophisticated regretting and rejoicing is, i.e., the more counterfactual outcomes agents consider, the more effort it will induce, simply because the derivatives of these intermediate cases will fall in between the derivatives for simple and sophisticated regret/rejoice.

Thus, we propose that sophisticated regret and rejoice could have evolved in environments inhabited by simple regretters, rejoicers, and homini œconomici since sophisticated forms of regret and rejoice push the agents to work even harder than any of the three previously considered types. Thus, in the remainder of this paper, we will analyze how sophisticated regret and rejoice compare to each other and under which ecological conditions should we expect one to be better than the other.

4. Motivating illustration: An evolutionary model

Let us apply sophisticated regret and rejoice in an evolutionary setting to illustrate their workings and their viability. To this end, we will need a handful of auxiliary assumptions about the concrete functional forms of subjective utility and cost. We will drop these assumptions again when we present our general results in Section 5.

Assume a well-mixed population of size unity with three types of individuals: homo œconomicus (HŒ), rejoicers (REJ), and regretters (REG). Their decision problem is the classical Robinson Crusoe economy problem of having to allocate their days' time between work, $x \in [0, L]$, and leisure, $L - x$. Working produces $\pi(x) = x$. We assume only $\pi(x)$ to be objectively given and relevant for eventual evolutionary success, i.e., evolution will simply favor those who work more. We normalize utility and cost of leisure to 0. Working comes with subjective effort costs of $c(x) = C \cdot x$ and can produce utility, $u(x)$, which is also subjective of course.

Critically, utility is where our agents differ: The first type, HŒs, has a canonical concave utility function, here: $u_{H\text{Œ}}(x) = u(x) = 2\sqrt{x}$. Their optimal effort, thus, is characterized by $u'_{H\text{Œ}}(x) = c'(x)$ and given as $x^*_{H\text{Œ}} = 1/C^2$. The second type, REGs, has the same baseline utility $u(x)$ but additionally regrets missed opportunities as described in Section 3. Their utility

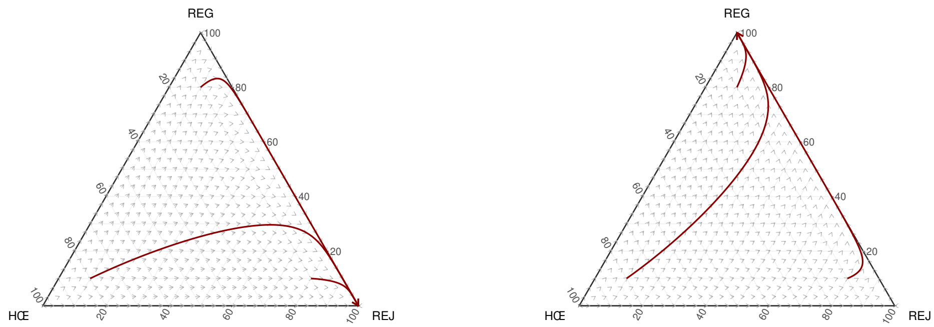

(a) $C = 2$, $L = 1$, $\phi = 0.1$

(b) $C = 2.5$, $L = 1$, $\phi = 0.1$

Figure (3) Two illustrating cases of evolutionary dynamics in populations of HŒs, REJs, REGs. The only difference between cases is in the cost parameter C.

is thus given by

$$u_{REG}(x) = u(x) - \phi \int_x^L u(t) - u(x)dt$$

Accordingly, their optimal effort is $x^*_{REG} = (4L\phi^2 + 4\sqrt{L}\phi + 1)/(C^2 + 4\phi C + 4\phi^2)$. Third, we have REJs in our population. Their utility is given by

$$u_{REJ}(x) = u(x) + \phi \int_0^x u(x) - u(t)dt.$$

as explained in Section 3 and their according optimal effort is $x^*_{REJ} = 1/(C^2 - 4\phi C + 4\phi^2)$ — as long as $2\phi - C < 0$.

With these three utility functions and respective optimal effort choices we are almost ready to study evolutionary dynamics. For simplicity, let these be described by the canonical replicator dynamics for well-mixed, infinite populations, i.e., let population change over time be described by

$$\frac{d}{dt}p_i = p_i[\pi(x_i^*) - \hat{\pi}(x^*)],$$

wherein p_i is a vector of the current population shares for types $i \in \{$HŒ, REG, REJ$\}$ and $\hat{\pi}(x^*) = p_{HŒ}\pi(x^*_{HŒ}) + p_{REG}\pi(x^*_{REG}) + p_{REJ}\pi(x^*_{REJ})$ is the current average population payoff. Note that these dynamics can be interpreted as either successful individuals' strategies—here, ways of feeling utility—being proportionately imitated more via social learning or simply as successful individuals having more offspring. We have no strict preference but stress that the biological (vs. cultural) interpretation of the dynamics is truly 'individualistic' in not requiring any capability to acquire social information on the individual's side.

Some systematic experimenting with the remaining free parameters, L, C, and ϕ indicates that for any positive ϕ, the HŒs are invaded and replaced by either REGs or REJs or, in rare cases, by a mix of the two. Figures 3a and 3b show two instructive cases.

Technicalities aside, what this illustration shows is that (i) personal standards induced by subjectively feeling either regret or rejoice can be instrumental and very effective in making agents (willing to) work more; and (ii) which of the two types (REGs or REJs) eventually works more depends crucially on the subjective cost structure of the task. In the following Section 5 we flesh out these results in more generality.

5. General results

Take the three types from Section 4, but let baseline utility, $u(x)$, now be quite generically any weakly concave function with $u(0) = 0, u' > 0, u'' \leq 0$. Furthermore, let (subjective) costs be described by a generic convex function, $c(x)$, with $c(0) = 0, c' > 0, c'' > 0$. Thus, agents are solving $\max_x u(x) - c(x)$ for optimal effort, for which the interior solution at x^* satisfies $u'(x^*) = c'(x^*)$. For HŒs utility is $u_{HŒ}(x) = u(x)$. For REGs utility is

$$u_{REG}(x) = u(x) - \phi \int_x^L u(t) - u(x) dt$$

and REJs have utility

$$u_{REJ}(x) = u(x) + \phi \int_0^x u(x) - u(t) dt.$$

Now, let us see how regretters and rejoicers perform in the Robinson Crusoe task. The derivative of u_{REG} is

$$u'_{REG}(x) = u'(x)[1 + \phi(L - x)]. \tag{1}$$

Notice that $\phi(L - x)$ is always positive except for the case $x = L$. Thus, REGs will always work more than HŒs. To see this, notice that $u'_{REG}(0) =$

$u'(0)[1+\phi L]$ and $u'_{REG}(L) = u'(L)$ and look at the left panel of Figure 4. One can see that the optimal effort choice x^*_{REG} by REGs is higher than x^*, the optimal effort choice of HŒs. Notice that this is a general result that holds for any assumed $c(x)$ and any assumed $u(x)$ with the given properties.

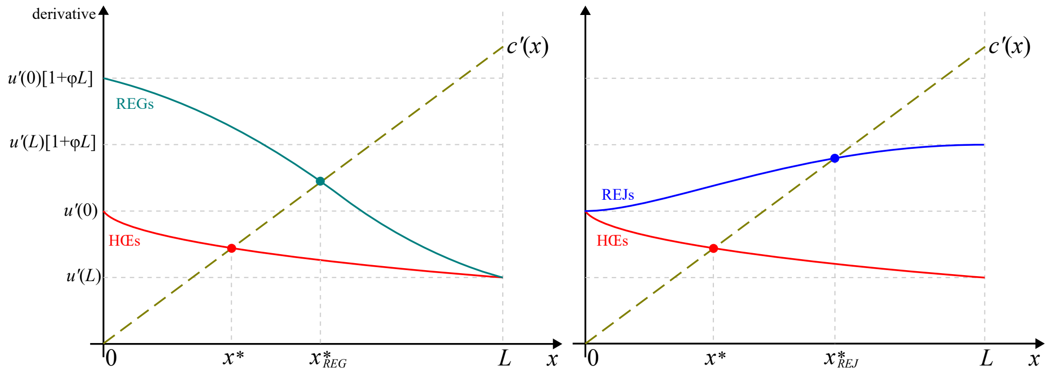

Figure (4) *Left.* HŒ and REG utility derivatives and the optimal effort choices $x^*_{REG} > x^*$. *Right.* HŒ and REJ utility derivatives and the optimal effort choices $x^*_{REJ} > x^*$.

Result 1. *For any concave utility and any convex costs, REG types put weakly more effort than HŒ types. They put the same highest effort when the cost derivative is low with* $c'(L) \leq u'(L)$*.*

Now let us carry out the same analysis for REJ types. For them the utility derivative is

$$u'_{REJ}(x) = u'(x)[1+\phi x].$$

Notice again that the term ϕx is always positive except for the case $x = 0$. Thus, the derivative for REJ types is always higher than for HŒ types. Figure 4, right panel, shows that $x^*_{REJ} > x^*$ in a similar way as with REG types. Again, this holds for any assumed $c(x)$ and $u(x)$.

Result 2. *For any concave utility and any convex costs, REJ types puts weakly more effort than HŒ types. They put the same highest effort when the cost derivative is low with* $c'(L) \leq u'(L)$*.*

These two results essentially mean that REG and REJ types will always invade populations of HŒ types, because these work less. Notice that this works for any positive ϕ. It is straightforward to see, moreover, that higher ϕ's make agents work even more than lower ϕ's. To see this, revisit Figure 4 and just shift the points $u'(0)[1+\phi L]$ and $u'(L)[1+\phi L]$ on the y-axis up a little (because ϕ grows); the green and blue curves then also shift up at one end, and the resulting x^*_{REG} or x^*_{REJ} are higher.

Result 3. *REG and REJ types with higher ϕ work more than a respective same-type agent with smaller ϕ.*

Thus, in any evolutionary process selecting for higher effort levels starting from a population of HŒs, once a mutation with very low ϕ appears that works more than HŒ types, selection will continuously favor mutants with higher ϕ's until some binding exogenous limit is reached.

Result 4. *The ϕ rises.*

So far, we have seen that both REGs and REJs perform better than HŒs in the Robinson Crusoe task. It is of course also interesting to pin down more precisely which 'ecological' conditions are more vs. less favorable for which of the two types. We can provide two general results regarding this question. We find that REG types can beat REJ types in specific situations: (i) in difficult environments with quickly-rising costs (high second cost derivative), and (ii) when maximum productivity exogenously increases (i.e., when L increases).

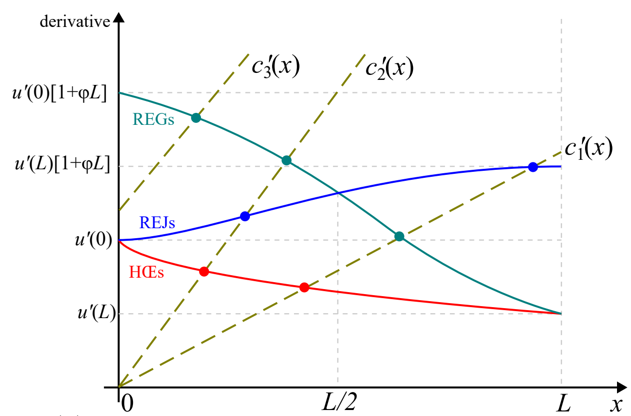

Figure (5) Optimal effort levels for different cost functions.

We start with the first one. Figure 5 shows three different cost derivatives: c_1', c_2', c_3'. Cost function c_1 rises slowly. Here, REG types put less effort than the REJ types (the projection of the green dot on the x-axis is further to the left than the projection of the blue dot). So, in case of slowly rising costs REJ types do better than REG types. However, when the cost is c_2 (the cost of putting zero effort is still zero, but the curve rises steeper than c_1) we have the opposite. Now REG types put more effort than the REJ types. Moreover, when the costs become even steeper and/or higher, like c_3, where the cost-derivative of zero effort is positive and above $u'(0)$, then HŒ and REJ types put zero effort, whereas REG types still put positive effort. This is significant, as its means that in c_3-type ecologies only REG types will work at all.

Result 5. *REG types prevail over REJ types in environments with steep (c_2) or high (c_3) costs.*

For our second result notice that the derivative u'_{REJ} does not depend on L, whereas the derivative u'_{REG} does. Their optimization depends on L and, from equation (1) above, it is trivial to see that when L rises, so does the derivative $u'_{REG}(x)$. Thus, the solution to

$$u'_{REG}(x) = u'(x)[1 + \phi(L - x)] = c'(x)$$

will move up as well. So, REG types with higher productivity work more than those with lower productivity. As a result, productivity increases will make REG types better off, because REJ types will not change their behavior in response while REG types will.

Result 6. *The REG types prevail over REJ types when productivity rises sufficiently (L increases) because they increase their effort while REJ types do not.*

To summarize. This analysis shows that the REG types beat REJ types in environments where costs are steep and/or when productivity increases sufficiently.

6. Discussion

Our analysis above has demonstrated how the psychological heuristics of sophisticated regret or rejoice could evolve in populations of homini œconomici and even fully replace them. Nevertheless, such sophisticated regretters and rejoicers are still 'selfish' in the sense that regret or rejoice over own utility does not immediately translate into social considerations.

In this section we outline an idea that, despite this, the very psychological mechanism that allows agents to take into account counterfactual outcomes and makes them more motivated (regret or rejoice) also can be used for cooperation and working together: it can be co-opted for social purposes. Suppose that in some social environment, agents—when computing their own regret or rejoice—also do the same for other agents involved. So, in each outcome x agent feels not only her own regret (a term like $-\int_x^L u(t) - u(x)dt$ above), but also regrets of other agents, which are similar utility terms that can be computed from the information about the values other agents receive in various outcomes. In this case, we can say that the agent feels *aggregated regret* at x for herself and all others involved. This aggregated value can be thought of as a measure of how 'socially desirable' outcome x is in comparison to other outcomes. If the sum of regrets of all agents in x is smaller than in some outcome y, then x is socially preferred to y. This logic of 'social' comparison of outcomes paves the way to the emergence of norms, personal and social ones, from the co-optation of individual regret or rejoice heuristics.

To see how this can be quantified, consider a game with N players and some set C of resulting allocations of utilities. Suppose that regret or rejoice that player i feels in outcome $x \in C$ is given by $r_i(x)$. This can be an integral as above or a sum of utility differences over a finite set C. Then for i, feeling others' rejoice or regret is the same as having a utility function of the form

$$u_i(x) + \phi \sum_{j \in N} r_j(x),$$

which adds regret or rejoice terms of all agents. Following [19] we can call $\eta(x) := \sum_{j \in N} r_j(x)$ the *norm function* that for each outcome x defines some number that is proportional to the overall social desirability of x deduced from how much overall regret or rejoice is felt in x. So, x is *more socially appropriate* than y whenever $\eta(x) > \eta(y)$.

Notice a very important property of this formulation. The only thing that was used to compute social appropriateness of outcomes was the information about payoffs in the game. As long as all players possess the same information about payoffs, they will all compute the same $\eta(x)$. This is important for two reasons. First, $\eta(x)$ is computed solely using the regret/rejoice mechanism that evolved for individual reasons and does not involve any specific giant-leap 'social' component. Second, $\eta(x)$ represents a *common belief* in social appropriateness of outcomes that emerges *endogenously* from separate, individual computations of each agent and not from their communication or observation of who does what. This idea shows that *social norms*—that presuppose common beliefs in social appropriateness of outcomes—can emerge from separate computations of individual agents who possess the same information about the game. And indeed, Kimbrough and Vostroknutov [19] show that this exact formulation where $\eta(x)$ represents sophisticated regret is doing remarkably well at explaining experimental results in social dilemmas, bargaining situations, and a wide variety of other contexts considered by behavioral economists.

We are aware that one may argue at this point that we did not really provide an account of how social norms emerge, but rather moved the difficult part of the explanation into the assumption that agents have the same information about the game, from which the same norm function is computed by each agent separately. This is indeed true: if agents have different information about the payoffs in the game then they might not agree on the social appropriateness of different outcomes. This property however may not necessarily present a problem for our argument, but may rather present another step to consider. We believe that in reality people do often disagree about social appropriateness of outcomes when having inconsistent beliefs about the payoffs. For example, part of the current climate change debate is rooted in the fact that people have different beliefs about the severity of the problem. This also holds for a plethora of other moral arguments which exist due to inconsistencies in factual beliefs. So, the problem in our view is not that people can never synchronize their beliefs about the world, they sometimes are very good at it; rather, the problem lies in understanding how cooperation and normative behavior can emerge in a world where beliefs are not exactly the same.

The idea that norms can be computed straightforwardly from information about a game actually allows us to shed some light on how exactly the various beliefs necessary for computations might end up being similar or different across many people. In fact, the problem of inconsistencies of beliefs can go deeper than beliefs about payoffs. It may be, for example, that players have different *social weights* attached to others when they compute $\eta(x)$. Specifically, it can be that the utility of player i is

$$u_i(x) + \phi \sum_{j \in N} \tau_{ij} r_j(x),$$

where $\tau_{ij} \in \mathbb{R}$ is a social weight that i puts on j (we can assume $\tau_{ii} = 1$). Players can have personal feelings towards each other, which determine these weights. For example, if player j has upset player i in the past, then τ_{ij} might be very low or even negative. In this formulation, we can call the function $\eta_i(x) := \sum_{j \in N} \tau_{ij} r_j(x)$ the *personal norm function* of player i. This formulation is similar to definitions in [20]; Kimbrough and Vostroknutov [19] also show how the idea of social weights explains behavior in social identity experiments by [21]. In principle, personal norm functions incorporate both payoff and social weight beliefs that can differ across players due to separate personal experiences.

We propose that, realistically, there are indeed a lot of differences and inconsistencies in beliefs across people and that normative views are actually different because of that. However, beliefs are also often *not extremely different* and norms based on inexact information can still help people to cooperate and work together. For example, value systems and social identities usually describe in detail who belongs to a social group (high social weight) and who does not (low social weight). The information about payoffs is also often discernible given that we all belong to the same species and have mostly similar needs. So, a group of agents who spent some time together and have roughly similar ideas about everyone's social weights and roughly similar ideas about payoffs will compute approximately same personal norm functions $\eta_i(x)$ that will allow them to cooperate to some extent—but likely not without some milder normative disagreements. The more similar the beliefs are in the group, the closer the situation will be to some commonly shared social norm function $\eta(x)$ and the better cooperation will become. This suggests that the process of evolution of norms as described here and the resulting cooperation can pick up even in situations where information about the world or attitudes of players towards each other are not exactly in sync.

Also notice that societies spend considerable efforts to make sure that beliefs in the population are synchronized. Rituals, traditions, religions, and education are all focused on creating common beliefs about payoffs (e.g., forbidden foods) or common beliefs about social weights (e.g., position in a status hierarchy). This suggests that the strict social norms we often see today could have evolved from much less coherent systems of beliefs by means of special cultural institutions that evolved in parallel to support synchronization of beliefs.

In summary, we have suggested that both personal and social norms could have evolved from individual psychological heuristics which were successful, because they can motivate agents—already in individual tasks—to exert more effort than others. Later, these mechanisms could have been co-opted to compute personal and social norms by simply adding regrets or rejoices of other agents to the same utility function. The idea that the same psychological mechanism may be used by everyone to compute norms suggests that common beliefs about social appropriateness of outcomes can arise without communication or observation of others. The mechanism also works when beliefs are not exactly the same, since the computed norms are continuous in the parameters and give similar results for sufficiently similar beliefs, thus producing intermediate levels of cooperation.

References

[1] M. F. Land, R. D. Fernald, The evolution of eyes, Annual Review of Neuroscience 15 (1992) 1–29. doi:10.1146/annurev.ne.15.030192.000245.

[2] S. Jessen, T. Grossmann, Unconscious discrimination of social cues from eye whites in infants, Proceedings of the National Academy of Sciences of the United States of America 111 (45) (2014) 16208–16213. doi:10.1073/pnas.1411333111.

[3] D. A. McLennan, The concept of co-option: Why evolution often looks miraculous, Evolution: Education and Outreach 1 (3) (2008) 247–258. doi:10.1007/s12052-008-0053-8.

[4] C. Bicchieri, The Grammar of Society: The Nature and Dynamics of Social Norms, Cambridge University Press, 2005.

[5] H. Sauer, The argument from agreement: How universal values undermine moral realism, Ratio 32 (4) (2019) 339–352. doi:10.1111/rati.12233.

[6] G. Lorini, Animal norms: An investigation of normativity in the non-human social world, Law, Culture and the Humanities 18 (3) (2022) 652–673. doi:10.1177/1743872118800008.

[7] H. P. Young, The evolution of social norms, Annual Review of Economics 7 (1) (2015) 359–387. doi:10.1146/annurev-economics-080614-115322.

[8] S. Legros, B. Cislaghi, Mapping the social-norms literature: An overview of reviews, Perspectives on Psychological Science 15 (1) (2020) 62–80. doi:10.1177/1745691619866455.

[9] S. Gavrilets, P. J. Richerson, Collective action and the evolution of social norm internalization, Proceedings of the National Academy of Sciences of the United States of America 114 (23) (2017) 6068–6073. doi:10.1073/pnas.1703857114.

[10] H. Rusch, The evolution of collaboration in symmetric 2×2-games with imperfect recognition of types, Games and Economic Behavior 114 (C) (2019) 118–127. doi:10.1016/j.geb.2018.12.005.

[11] J. Newton, Evolutionary game theory: A renaissance, Games 9 (2) (2018) 31. doi:10.3390/g9020031.

[12] S. Gavrilets, Coevolution of actions, personal norms and beliefs about others in social dilemmas, Evolutionary Human Sciences 3 (2021). doi:10.1017/ehs.2021.40.

[13] V. Calabuig, G. Olcina, F. Panebianco, The dynamics of personal norms and the determinants of cultural homogeneity, Rationality and Society 29 (3) (2017) 322–354. doi:10.1177/1043463117717233.

[14] G. Loomes, R. Sugden, Regret theory: An alternative theory of rational choice under uncertainty, The economic journal 92 (368) (1982) 805–824.

[15] G. Coricelli, H. D. Critchley, M. Joffily, J. P. O'Doherty, A. Sirigu, R. J. Dolan, Regret and its avoidance: a neuroimaging study of choice behavior, Nature Neuroscience 8 (9) (2005) 1255–1262.

[16] M. Fioretti, A. Vostroknutov, G. Coricelli, Dynamic regret avoidance, American Economic Journal: Microeconomics 14 (1) (2022) 70–93.

[17] N. Camille, G. Coricelli, J. Sallet, P. Pradat-Diehl, J.-R. Duhamel, A. Sirigu, The involvement of the orbitofrontal cortex in the experience of regret, Science 304 (5674) (2004) 1167–1170.

[18] B. Timberlake, G. Coricelli, N. Bault, Regret, responsibility, and the brain, in: The Moral Psychology of Regret, Rowman & Littlefield Publishers, 2019, p. 165.

[19] E. Kimbrough, A. Vostroknutov, A theory of injunctive norms, mimeo, Chapman University and Maastricht University (2023).

[20] Z. Bašić, E. Verrina, Personal norms—and not only social norms—shape economic behavior, MPI Collective Goods Discussion Paper (2020/25) (2021).

[21] Y. Chen, S. X. Li, Group identity and social preferences, American Economic Review 99 (1) (2009) 431–57.