

ON THE NORMATIVE RELEVANCE OF SYSTEMIC RISK

PIETRO MAFFETTONE

Dipartimento di Scienze Politiche

Università di Napoli Federico II (Italy)

pietro.maffettone@unina.it

RYAN MULDOON

Department of Philosophy

University at Buffalo (USA)

rmuldoon@buffalo.edu

ABSTRACT

The social world is permeated by risk exposure, and especially systemic risk, that is, risk we cannot really hedge against or protect ourselves from. Systemic risk is determined by the basic structure of a social system and affects the kind of choices we are able to make in our lives. We argue that when systemic risk is ‘too low’ society becomes stagnant as it does not allow for processes of creative destruction that, according to a long tradition of economic thinking, are at the core of what allows for growth, and thus progress. In the same way, when levels of systemic risk are too low, the range of option risks that individuals can decide to bear is itself too low and thus hampers their self-respect. At the same time, we will argue that when levels of systemic risk are too high, society runs the risk of marginalizing the potential contributions to innovation and growth of a large part of its members, for when there is too much systemic risk, too much of one’s life is uncertain, and thus investing in one’s future becomes less important. Excessive levels of systemic risk entail a lesser ability to pursue one’s conception of the good.

KEYWORDS

Risk, systemic risk, distributive justice, classical liberalism, high liberalism.

1. INTRODUCTION

Classical liberals and high liberals dispute over the appropriate size and function of the state, and how that shapes and scopes individual agency.¹ Classical liberals

¹ For an account of the distinction see Freeman (2011). According to Freeman, classical liberalism gives greater weight to economic liberties compared to high liberalism. In his words: “Where liberals primarily disagree is on the nature and status of economic rights and liberties, including the extent of

favor leaving the state small, and granting a wide scope to individual action. As much as possible, the state should leave individuals alone to pursue their life plans. High liberals worry that “leaving individuals alone” means having them walk a tightrope without a safety net, and so favor shoring up that safety net by expanding the scope of state action (Rawls, 2001). Classical liberals see the expansion of the state as a threat to individual choicemaking, and a loss of autonomy (see Tomasi, 2012). They are right. High liberals see a more minimal state as incapable of ensuring that differently-situated individuals really have equal rights, and “freedom” to sleep under a bridge is no freedom at all. They are also right (see O’Neill and Williamson, 2012).

Rather than attempt to settle the dispute between classical and high liberals, we wish to situate their debate in terms of how one sets an appropriate range of allowable systemic risks in society. Both disputants are very good at demonstrating challenges with the excesses of the other, but can sometimes be blind to their own shortcomings. We ultimately argue in favor of what we call ‘interval limitarianism’, which suggests that there is a limited range of levels of systemic risks that individuals can be safely exposed to.² Too much risk, and agency is undermined because it is impossible to establish firm footing. Too little risk means agency is undermined because the exercise of choice has been removed from too many facets of one’s life. By reorienting discussion around a risk framework, we can more fully appreciate the contributions of classical and high liberalism, and the conditions under which liberal autonomy can flourish. To see the benefits of this approach, let us first think about risk.

Risk occupies a relatively unique position in our lives (Hansson, 2018; 2013). Without a willingness to take on some risk, society would stagnate. No risk means no entrepreneurship, no political change, no innovation in art, and no social and economic mobility. Risk is a necessary component of the economic, political, social and artistic dynamism that makes social progress possible. But too much risk, and we have a different form of stagnation (see Heath, 2001). When there is too much risk, individuals can’t reasonably plan for their futures. Companies can’t invest. Too much risk forces us into a debilitating form of presentism and conservatism. In this extreme, the best we can do is scramble to try and maintain the status quo. Between these two extremes, however, we can find levels of risk that are conducive to productive forms of dynamism without overwhelming our ability to plan for our futures.

freedom of contract and rights to private property in land, raw materials, and other productive resources. Classical liberals generally hold that the economic liberties are to be regarded and guaranteed as among the basic liberties; or if they are not strictly basic liberties, then economic liberties resemble basic liberties in that they can only be restricted for special reasons” (Freeman, 2011: 20)

² We borrow the term ‘limitarianism’ from Robeyns (2017a), yet we use it in a different way as the qualifier ‘interval’ clearly suggests.

We can think of risk as being exposed to a gamble (for different approaches to risk in economics see Barr, 2012; Landes, 2015). Not all risks are equal, of course. Some have clearly defined odds, where others produce uncertainty.³ Some risks are relatively balanced between upside and downside, others are skewed in one direction or another. While these distinctions are undoubtedly important, unless we specify otherwise, here we use risk as a decrease of certainty over outcomes, where those outcomes could be positive or negative. One important distinction that we will rely on is between what we will call systemic risk and option risk. Systemic risk stem from the basic structure of society. Given a particular basic structure, individuals cannot easily opt out of being exposed to this form of risk. Option risk, on the other hand, is to a large extent up to individual choice.

While one cannot, for instance, materially change her odds of whether there will be a recession when she is looking for a job, she can to a much larger extent choose a career that is more or less risky in terms of the opportunities that are likely to be available to her in the future.

As intuitive as this simplified picture is, it suggests that systemic risk is an odd kind of social phenomenon. It is a social good when there's roughly the right amount of it, but too much of it, and it becomes a social bad. This is unlike most other 'social goods' or 'social bads'. For instance, while wealth and income may have diminishing marginal benefits for individuals, there isn't a level of wealth and income when more of it becomes a burden rather than a benefit. Instead, systemic risk looks more like medicine – in the right dose, it is extremely helpful, but if there's too little it can't work, and if there's too much of it and it becomes a poison. We will say more on this below.

Our main claim in this paper is to argue that an important way to morally assess a social system from the perspective of its effects on individuals is through the prism offered by the idea of systemic risk. The reason for the latter conclusion is that the amount of systemic risk in society affects the kind of lives persons can lead. It does so by altering their ability to make autonomous choices, and by changing what kind of future these persons can collectively aspire to.

To make our position explicit, we argue that for both individuals and for society, systemic risk exposure is a key factor to innovation and growth. There is not a way of creating a satisfying systemic risk-free society. Our lives fundamentally involve uncertainty, and over-aggressive efforts at eliminating sources of systemic risk would strip out much that we have reason to value as individuals and as members of a given society. For it would deprive individuals from making important choices, and

³ The reader may note that we are conflating risk and uncertainty. While we recognize the difference between the two concepts, our argument does not hinge on this difference, so we combine the two concepts for simplicity.

society from being able to rely on incentives to innovate and make progress. However, when one is exposed to too much systemic risk, it is not merely bad because there are more gambles to go sour, but rather because when too many aspects of our lives are exposed to significant risks, our agency is undermined. When we have no firm footing because everything is up for grabs, we can't hope to make reasonable choices. The same applies to a society at large. Too much systemic risk, at the limit, might hamper the willingness of individuals to make risky choices, to be entrepreneurial, to try out new ideas etc. Recent work in behavioral economics suggests that our ability to effectively evaluate choices is a function of the number of high-stakes choices we make – our “cognitive bandwidth” is limited (Mullainathan and Shafir, 2013). The more we have to evaluate choices about the basics of our lives, the less we are equipped to effectively evaluate future-looking choices.

Put differently, we shall argue that when systemic risk is ‘too low’ society becomes stagnant as it does not allow for processes of creative destruction that, according to a long tradition of economic thinking, are at the core of what allows for growth, and thus progress (see Schumpeter, 1994; Cowen, 2017). In the same way, when levels of systemic risk are too low, the range of option risks that individuals can decide to bear is itself too low and thus hampers their self-respect; individuals’ lives become the product of wider social choices rather than, at least in part, being the outcome of their efforts and decisions. At the same time, we will argue that when levels of systemic risk are too high, society runs the risk of marginalizing the potential contributions to innovation and growth of a large part of its members, for when there is too much systemic risk, too much of one’s life is uncertain, and thus investing in one’s future becomes less important (see Hacker, 2006). Similarly, when too much systemic risk characterizes a social system the lives of individual citizens become less autonomous. Excessive levels of systemic risk entail a lesser ability to pursue one’s conception of the good. In this picture, citizens’ autonomy is degraded and their ability to decide and pursue what they take to be of value over the course of a lifetime is deeply affected.

2. RISK, OPTION RISK, AND SYSTEMIC RISK

Risk is not an easy concept to pin down. The concept itself risk is not univocally defined (Hansson, 2013). Both in everyday language and in the literature on risk, authors often use the concept in different ways. In cost-benefit analysis, risk is defined as the probability that an event will take place multiplied by the (negative) utility associated to the materialization of the event. Risk is usually distinguished both from certainty, where the probability that an event will take place is 1, and uncertainty, where the probability that an event will take place is unknown.

The crucial question, for our purposes, concerns the way in which we can square the normally accepted definition of risk with the also commonly accepted idea that

a life without risks would be, coarsely described, a dull one. The problem is clear enough. We usually define risk in the following way:

$$R = P(e) \cdot U(e)$$

Where $\{U(e) < 0\}$ is the utility associated to a given event taking place.

And $\{0 < P(e) < 1\}$ is the probability of the event taking place.

In this picture, risk always refers to the potential occurrence of an event with which we associate a negative utility. Thus, strictly speaking, the main reaction we should have towards risk would be to try to minimize it. However, clearly, the normative judgment about risk minimization is, at best, incomplete. Why so? Because most risks are associated with choices that we deem worthwhile, and that have potentially very beneficial outcomes attached to them (see Adams, 1995). One way to informally capture this idea is that, for each risk, there is an opportunity. Minimizing risk is, then, desirable only ‘other things being equal’ – which they rarely are.⁴

Let’s go through some basic illustrations. Every time you exit your home, you incur some risk of injury that would not materialize were you to stay at home. Even assuming that, for the sake of argument, nothing can ever happen to you if you stay at home, what exactly are we to conclude from the fact that not going out would allow you to minimize risk of injury? Most people would say: not much at all. Now, consider the example of a medical procedure. The procedure will inevitably force you to incur some form of risk if you consent to undertake it. Yet, presumably, there would also be clear benefits attached to it (there are exceptions, but let’s leave those aside). Clearly, if risk minimization is the only goal, then one should not undertake the procedure. Yet, also clearly enough, the benefits connected to the procedure would be lost as well. Thus, if it is rational to maximize expected utility, both the potential benefits and risks associated to a choice need to be taken into account.

As we have said in the introduction, a further important distinction about the nature of risk is the one between option risk and systemic risk. Most analyses of risk tend to depict the latter, as we have seen, as a description of a particular situation faced by a specific individual at a specific point in time. However, this might suggest, even if just implicitly so, that risk(s) are something that individuals necessarily face as decision makers – that they can decide whether they should take them or not. Yet, as much of the recent literature on the ethics of risk has shown (see for example Frick, 2015; Lehman, 2008), some of the most interesting ethical aspects connected to risk are generated precisely by the fact that those that bear the risk and those that

⁴ Note that this is the reason, within the framework of cost-benefit analysis, that the main goal is not risk-minimization, but the assessment of the balance between risks and benefits that a given action, project, policy etc. might carry with it (for an overview on cost-benefit analysis see Adler and Posner, 2006)

decide that a certain risk should be borne are not one and the same agent (see Hermansson and Hansson, 2007).

To illustrate, we often assume an interactional picture of risk: whether chosen or imposed upon them, risks are faced by agents as isolated events. To picture the risks faced by most individuals as separate instances of risk may invite the idea that any risk can be the object of a free choice. Yet, that is potentially misleading. Many of the risks people are exposed to are not the result of the choices they make. This is, roughly, the kind of idea we have in mind in this paper when we speak of systemic risk: risk that is, for the most part, unchosen. Systemic risk is not a feature of the decisions made by individuals, but rather a characteristic of the kind of social environment in which they happen to be placed; it defines the kinds of options they have access to rather than the choice of a specific option within one's option set.

Examples of these kinds of risk are, in our view, ubiquitous. To provide just two simple ones: depending on the kind of economic system one happens to operate within, the risk of not having enough economic resources in retirement might, or might not be, something individuals are exposed to: in a command economy with fully guaranteed state pensions the risk is absent, in a market economy in which welfare entitlements are non-existent the risk is clearly there.⁵ Similar remarks would apply to health care. The risk of being unable to pay for one's medical expenses is likely to be heavily dependent on structural features of a social system. For, if well-functioning free public universal health care is in place, chances will tend to zero, while they would be clearly higher, other things being equal, in a privately run system in which access to care depends on ability to pay.

Of course, to say that in a free market economy we are constitutively or systemically exposed to pension risk (or to what we can define as health care risk) is not to say that all individuals are exposed to the same kind of probabilities of not having enough income in retirement. After all, different people will have different incomes and savings rates that will influence their retirement income (and, in the same way, their ability to pay for medical expenses). More broadly, systemic risk can be more salient for some than for others. Regardless, we want to maintain that there is a sense in which this kind of risk (i.e. systemic risk) is still relevant as a feature of the *system* in question. What shall say more about this specific point in the next part of the paper.

What is the connection between systemic risk and choice? It might be thought that looking at a social system from the perspective of systemic risk defeats the very

⁵ While the command economy might not expose anyone to what we can call pension risk in a direct way, there might be other sources of risk that bear on the case, for example, that the state might become unable to pay such entitlements or that these might gradually lose their value. The latter might be defined as an element of sovereign risk or counterparty risk, where the relevant counterparty is the government or state. This is not relevant to us here and we can simply add to the example the guarantee that state pensions are, *arguendo*, regularly paid and stable in value.

purpose of our normative analysis. It might be thought, in other words, that to look at risks that are largely unchosen does not leave much space for agential control. That is not the case. On our framework, systemic risk largely is a constitutive feature of the option set that individuals face. Within this set, and given the features of characteristics of this set, individuals make choices and exercise their agential capacities. Looking at the set is, nonetheless, very important because the nature of the set itself will tell us something normatively significant about the nature of the choices that individuals are making. As Sen (1999; 2009) has illustrated, we have reason to be interested in the real choices that individuals can make, and whether those are adequate for pursuing a life of one's choosing.

Put differently, systemic risk can be understood as a proxy for the kind of choice environment that individuals will face in social life. The kinds of choices they are required to make for themselves and for which they will bear the negative effects should they materialize. Systemic risk, in other words, is the social, economic and political ecosystem individuals need to interact with but cannot necessarily shape in the short to medium term. The question we need to ask, then, is what kinds of features should this ecosystem have for the lives of democratic citizens to go well? Put another way, what choice environments allow them to be able to lead lives in which they can both be autonomous agents without needing to scramble for safety at every turn, and yet make sure that, collectively, they live in a dynamic society that would allow them to receive the benefits of high real growth and technological change.

Finally, note that it is here that we see an important connection between our understanding of systemic risk and Sen's conception of capabilities (for a recent overview see Robeyns, 2017b). Less systemic risk entails more control over one's option set. Systemic risk minimization strategies work by making the presence of mandatory risky choices less prominent within one's complete option set, yet it does so at the price of cleaving off portions of one's option set altogether. However if collectively adopted, systemic risk minimization strategies would severely affect the aggregate outcomes that a society can aspire to. It would make for a poorer, less interesting and thriving social world. In addition, systemic risk minimization strategies can be costly, at least when done by a third party (i.e. being the outcome of an institutional decision), as they might affect one's ability to shape her life. This is central to Sen's conception of the capabilities approach. A wider option set doesn't assure the best of all possible worlds if we do not know how 'easy' or 'available' are each options within the set, yet, reducing the set itself (by reducing systemic risk) cannot be something we should accept lightly, for it might very well have significant costs, morally and materially.⁶

⁶ For a more formal depiction of the connection between risk and capabilities, we propose that the object under inquiry is the portfolio E of potential outcomes in some social context. For each

3. ADOPTING A SPECIFIC (SOCIAL) POINT OF VIEW

Following from the previous section we can state that what matters, in the end, is the level of systemic risk that exists in a given social system. Yet, at the same time, we should also note, as we have already alluded to in part II, the kind of social position from which we should be inclined judge the amount of systemic risk that is present in any given society will be of great importance when it comes to its overall acceptability. One intuitive rationale for distinguishing between different social groups is that, as we have seen above, these groups have different endowments which can make managing systemic risk more or less burdensome to them. To illustrate, high-interest, short-term loans like payday loans that can trap low-income people in penury may be highly advantageous to a house flipper or a stock market trader. Those loans allow people who can readily convert capital into profits to make far more money. Thus, we can legitimately ask ‘where from’, so to speak, should we consider how much systemic risk we should recommend for a society to accept?

The choice of which socially representative group one should adopt to judge acceptable levels of systemic risk in society is of course a difficult one. Here, in keeping with much of the so-called ‘high liberal tradition’, we are inclined to look at systemic risk from the perspective of the least well off individuals. Adopting the perspective of the least well off is of course indeterminate unless some form of definition is also specified. The basic idea is that the least well off representative individual is the one for which the aggregate potential negative effects of systemic risks are likely to be higher over the course of a complete life. Clearly enough, this kind of definition might map quite well onto other kinds of ways we tend to use to describe the least well off members of society (see Rawls, 1999: 83-84), for the ability to manage systemic risk (as we have defined the idea) will be influenced by, among other things, one’s level of income, wealth, job security (and thus skills), etc. To see why, the reader can return to the examples of systemic risk we have offered a few

outcome e in E , there a utility $U(e)$, and a probability $P(e)$. A risk minimization strategy is defined as a portfolio-trimming strategy, where $\text{Card}(E)$ is reduced. We take it as a premise that portfolio trimming will be a maximin strategy that reduces the overall expected utility of the portfolio. (Were this not the case, we would be merely eliminating irrelevant alternatives). From a baseline portfolio, we can see that expanding the portfolio (adding to the cardinality of E) can happen in three ways: we can decrease expected utility, be neutral with respect to expected utility, or increase expected utility. Finally, we can compare portfolios E and F not just by their cardinality and expected utility, but by their variance in outcome utilities. When $EU(E) = EU(F)$, risk preferences determine whether E or F is favored. On our framework, we can consider systemic risk to be the distribution of portfolios across representative agents. These represent the choice sets available to the agents. Option risk is then captured by the choices that agents make (or gambles that they accept) given the portfolios they receive that fix their menu of choices.

paragraphs above. When it comes to not having enough income in retirement and not being able to access health care, clearly, income and wealth might be good proxies for a person's ability to manage systemic risk.⁷

However, defining a social position does not justify why we should look at the world through it. Thus, what isn't entirely clear is why exactly one should focus on, insofar as systemic risk is concerned, the least well off members of society. At least since the publication of Rawls' *A Theory of Justice* (1999), we are accustomed to believe that, morally speaking, the interests of the least well off should have some kind of added weight in our deliberations about how to structure a political community (see Parfit, 2002). Here, we do not wish to deny the latter claim, but would like to add to it by proposing a different kind of rationale for concentrating on the least well off when we are dealing with levels of systemic risk.

One of the main normative questions for any social system is, on our account, how much systemic risk there should be – that is, how we shape the risks that citizens are necessarily exposed to in virtue of their being citizens in a given social position. Being in a market economy, for instance, forces individuals to make certain choices that wouldn't exist were they in a planned economy. But though we are primarily concerned with systemic risk, the very reason for caring about systemic risk is that these choices will interact with the option risks that people take on. For instance, if one is closer to the poverty line, an insufficient social safety net may result in someone taking on payday loans or auto title loans to cover liquidity problems. Or perhaps the person may take on a job with "just in time" scheduling, which can make it difficult to pre-arrange childcare. These are option risks that the person takes on, but would plausibly never choose had they been exposed to a different portfolio of options shaped by the systemic risk level that they encounter. And, we would like to claim, these kinds of choices will be disproportionately forced upon the least well off members of society, and that this kind of phenomenon is likely to worsen over time.

To see why, consider the following. Our sense is that there might be something akin to a gravitational dynamic in play. The gravitational dynamic can be described in the following way. First, fix a given level of systemic risk in society. By definition, in our account, the least well off members of society are those who find it more

⁷ To say that income and wealth will often track exposure to systemic risk is not, however, to say that they are identical. In other words, income and wealth will often, but not always, be good proxies for vulnerability to systemic risk. To see this, consider the vulnerability to systemic risk of an economically poor young adult who just graduated from a good university with a large amount of debt. Now compare the latter to the vulnerability to systemic risk of an older person who has, ex hypothesis, savings as opposed to debt, but suffers from a chronic medical condition that is very costly to treat and not covered by his/her health insurance. The first person is economically poorer than the second, yet it is the second that seems intuitively more vulnerable to systemic risk compared to the first.

difficult to manage such risk. Then, consider what happens when you vary the level of systemic risk in society. Imagine, more specifically, that you increase it by a given delta. What would happen to the least well off in this new configuration? Our intuition is that their ability to manage systemic risk would be more than proportionally affected by the increase. Put differently, there is a sense that, the less one is able to manage systemic risk, the more one becomes disproportionately vulnerable to its increase. The potential negative effects of increased systemic risk may (disproportionately) flow, so to speak, to those who are already most vulnerable to them.

This is, coarsely stated, the underlying intuition we have. How can we make sense of it? Consider here a very simple model of political economy. In each round, questions of systemic risk (i.e. its level in society) are put to a vote. People can vote their preference or not vote at all. Whatever policy that receives the most votes is put into effect, and the level of systemic risk is updated. One might think that this would ideally set the correct amount of systemic risk in a social system: as different kinds of citizens experience different levels of systemic risk, they can vote their interests, and the system will adjust itself accordingly (modulo some background constraints on what's possible). We would not only get what was a more or less ideal level of systemic risk, but it would also come with significant normative justification in the form of citizen consent. If voting was somehow compulsory or was "free" in the sense that no one took on any costs to vote, then we should expect something like this initial assumption to come true. However, if we instead more realistically assume that there are costs associated with voting, this story changes.

Let's assume that one's likelihood of voting is at least partially a function of their vulnerability to the effects of systemic risk that they are currently exposed to. That is, the more they are making choices brought on by volatility of future outcomes or by bad luck from previous outcomes, rather than being able to weather some variance, the less likely they are to vote. This is a natural assumption: the more one is focused on their daily survival, the less they are going to plan or try to shape their future. If we assume this kind of voting pattern, then we should expect non-voters to get exposed to ever-more risks (perhaps up until a saturation point). This doesn't need to be out of malice by anyone else, but rather because insofar as voters are weighing tradeoffs between interests, it becomes free to impose costs on people who don't vote. Indeed, if there is a saturation point, we may find that the levels of systemic risk adopted through voting may be continually pushed 'up', and, as a consequence, its downstream effects would continue to increase and become less and less manageable for those who are worse and thus who find participation to be costly. This basic account is going to share a great deal in common with a preferential attachment model. The gains from any given level of systemic risk will shift upwards, while the burdens will shift downwards. This doesn't require malice or a conspiracy. Just differential burdens becoming differential voice, and that difference in voice shaping future burdens.

In this picture judging a social system, in terms of systemic risk, from the perspective of the least well off does not require the (morally sound) prioritarian assumption that the interests of the least well off are to count for more. Rather, the attachment model of political economy we have just sketched suggests that, over time, the least well off will become disproportionately affected by every marginal increase in systemic risk. They will thus reach what we can call the ‘systemic risk frontier’, that is, the level at which their contribution to a dynamic society and their ability to make autonomous choices are both undermined, much sooner than other groups in the political community. Yet, if society wants to show equal respect and concern to each and every one of its members and if it wants to draw on the productive contribution of all of them, then, no substantial parts of it can be allowed to reach such systemic risk frontier. And in turn, this provides reasons to concentrate on the least well off.

4. SYSTEMIC RISK: A DIFFERENT WAY TO LOOK AT SOCIETY?

In part I of the paper we have elaborated on the distinction between option risk and systemic risk. In part II, we have tried to justify a specific point of view from which judgments about acceptable levels of systemic risk in society should be made. Here, we shall try to say something more about the nature of systemic risk, namely, that we believe the latter idea to offer an important and underexplored prism to look at different kinds of social worlds. In what follows, we shall not claim that systemic risk is itself a distributive currency. Nonetheless, it is important to note that, much like currencies of distributive justice, systemic risk allows us to capture something relevant about the kind of society we live in.

Distributive currencies usually play specific roles within accounts of distributive justice. For example, they allow us to establish the correct perspective according to which we are entitled to make comparative judgments about who is faring better or worse. In addition, it can be argued that distributive currencies indicate what we ought to consider of value from a public perspective. Thus, putting the latter two observations together, we can say that, at the very least, distributive currencies give us a way to compare how people are doing from the perspective of what we believe we ought to consider the publicly relevant kind of standard.

Leaving several complications to one side, we can also say that the amount of systemic risk in a social system, especially when it comes to those that are most exposed to it, can play a similar role. Put differently, systemic risk can allow us to capture something we ought to consider of some value from a public perspective. It does so, we want to suggest, because it enables us to acquire a different perspective on the kinds of lives individuals (and especially those who are more exposed to systemic risk; i.e. what we have defined as the least well off) are allowed to live. Systemic risk gives us an interpretive tool to compare how different people’s lives

go, and it does so from a perspective that is publicly relevant, at least insofar as different levels of systemic risk are the result of collective choices about the basic structure of society. Thus, while we shall not argue that systemic risk is itself a currency of distributive justice, we can say that it can, for limited purposes, operate in a similar way.

With the aforementioned caveats and clarifications in place, we proceed by situating systemic risk in the wider literature on distributive justice by noting some of the main differences between systemic risk and the major distributive currencies. Our analysis is, needless to say, highly selective in that we will concentrate on features of distributive currencies that will allow us to bring out what we think is most distinctive about systemic risk.

To start-off, consider the fact that, roughly speaking, most theories of distributive justice are concerned with the distribution of ‘social goods’ broadly defined. We can think of income and wealth, and various forms of opportunities and entitlements as proxies.⁸ Two of the shared features of these otherwise very different currencies of justice are that they all purport to confer some kind of advantage to recipients and, perhaps more controversially, they all seem subject to diminishing marginal returns. Restricting oneself to income and wealth for the sake of simplicity, it seems relatively clear that for any individual, any additional ‘unit’ of income, and/or wealth, will represent some form of added advantage and at the same time that the advantage that is conferred by the n th unit of income, and/or wealth, will be less than the advantage conferred by the n th-1 unit. Yet, it is also clear that, bar specific exceptions, there will be no point at which an extra unit of income, and/or wealth, will have a negative effect on the recipient.

While less amply discussed, we can also think about the kind of features that would be displayed by ‘social bads’ if were to adopt them as a currency of distributive justice. Here think of, for example, pollution, or lack of access to medical care, or whatever currency that we can conceive to have a negative impact on an individual. While social goods seem to be subject to decreasing marginal returns, we can instead picture social bads to be subject to negatively increasing marginal returns. In addition, we can also conceive some social bads to be subject to what we can call threshold effects, namely, that whatever the kind of disadvantage that is conferred by the social bads in question, there exist some specific point after which the quantity of disadvantage becomes close to infinite. Pollution provides an easy illustration of these dynamics. A slightly polluted environment to live in is clearly a social bad

⁸ Examples are Rawls’ broader resourcist metric of primary goods (1999), Arneson’s opportunity welfarism, Cohen’s luck-based view centered on access to advantage (1990), or Sen and Nussbaum’s focus on capabilities. Note also that there is a clear sense that our discussion in this part of the paper fits better within a resourcist framework for one of the features of resourcist currencies is that they imperfectly track advantage to real individuals placed in a given social system and given their real endowments.

that confers some kind of disadvantage on individuals. As the level of pollution grows, the effect that latter has on those who will experience it will, generally speaking, become proportionally larger. At some point, however, increasing (negative) marginal returns will not fully capture the effects of pollution on those who experience it simply because the individuals in questions might find the environment in which they live not hospitable to human life anymore and at the limit, they might simply die from the pollution.

Put the aforementioned observations together, we can say that, whether one distributes social goods or social bads, one can be relatively safe in the knowledge that the amount of those goods that is to be distributed will not, by itself, affect the (qualitative) nature of the effects on the recipients. Using more technical language, we can say that social goods and bads can be represented by classes of functions $Y = F(q)$ relating quantity of the currency and effects on recipients with first derivatives which, at least most of the time⁹, do not change sign.

How can systemic risk be understood in terms of the main features we have so far highlighted? In other words, what are the effects of systemic risk on those who experience it? As we have argued above, systemic risk requires us to adopt a much more complex outlook on a social system. Life without it can be dull and might deprive individuals of important experiential and expressive elements connected to their agential capacities. Even more importantly, a society that severely curtailed systemic risk would simply not be able to attain the kinds of results that a more dynamic society characterized by a more dynamic environment can aspire to. However, too much systemic risk and the only thing that is left for individuals to do, especially for worse off individuals, is to scramble for safety. Thus one initial and important feature of systemic risk from the perspective of the evaluation of a social system is that, unlike most distributive currencies, it cannot easily be depicted as either a social good or a social bad. The nature of its effect on those who are exposed to it will simply vary according to the level they are exposed to.

Our claim here is not about where or how to locate the point at which additional systemic risk changes the kind of effect on the exposed. Depending on how one decides to conceptualize the situation (for example, in empirical terms by asking real persons? Or in philosophical terms by deciding what kinds of attitudes idealized citizens of a political community should have towards risk?),¹⁰ the point of inflection for the effects of systemic risk on individuals will vary. Our claim is instead

⁹ This might not hold when we consider threshold effects.

¹⁰ More fine-grained distinctions are also possible. Leaving aside systemic risk for a second, consider how much option risk different persons are willing to bear. If we wanted to assure that everyone had the level of risk they were most comfortable with, then we would have reason to deviate from a flat distribution of risk, as some people are more risk-loving, and some are more risk-averse. Here, see Cohen and Einav (2007) for an empirical estimate of the distribution of risk preferences. They

more theoretical in nature and pertains to the basic idea that the effects of systemic risk on individuals can be depicted as having an ‘inverted U’ shape – starting at very low levels of exposure, additional units will contribute to a person’s life and her society’s improving, yet as the quantity of systemic risk one is exposed to grows, there will be a point at which any additional increment will have negative effects on a person’s life, and on her society as well.

5. SYSTEMIC RISK AND INTERVAL LIMITARIANISM

Let’s recap what we have argued so far. First, we have introduced a distinction between systemic and option risk. Systemic risk characterizes the option set from which individuals can deploy their agential capacities. Option risk has been depicted as the kinds of choices, in terms of which risks and/or how much risk, individuals can make given their option set. Furthermore, we have analyzed the nature of systemic risk as possessing distinctive features at least insofar as we compare it to other well-known social goods or social bads. As we have seen, systemic risk offers a special way to look at how social system is structured and how individuals may fare within it, for, the amount of such risk we witness in a given social system will have to stay within some kind of interval for individuals and society to benefit from it. The basic idea is that, as we have stated in the introduction, much like medicine, we need some, but not too much, systemic risk in a society. Life without being exposed to systemic risk would be unbearable and yet when too much systemic risk characterizes a society, life becomes a struggle rather than a fulfilling attempt to engage with one’s goals.

In what follows, we try to tie together these arguments into a more unified whole, and will concentrate on the relationship between systemic risk and individual agency and autonomy. More specifically, we will try to better explain our intuitive view of the role systemic risk plays in persons’ experience by offering a clearer justification for the fact that people need neither too little nor too much of it. Furthermore, we will try to offer a conceptualization of the relationship between systemic and option risk that is grounded in this kind of justification.

The upshot of our argument is that, when systemic risk is concerned, we should adopt a view we have labelled, in the introduction to this paper, ‘interval limitarianism’. More precisely defined, interval limitarianism states that one (and yet clearly not the only or main) feature of any social system or basic structure of society is the

find systemic differences across age, gender, wealth and income. If we ignored preferences, and instead looked at other criteria, we might think that we may want to pay attention to the effects of the risk, and look to be more equitable in the distribution of the effects of risk, rather than just the risk itself. For instance, wealthier individuals are better able to benefit from financial risks than poor individuals, simply because they are better able to place more bets given their greater resources.

way in which its members are exposed to systemic risk (i.e. the level of exposure). Such exposure to systemic risk must lie within a specific interval. The interval sets the boundaries of the acceptable amounts of systemic risk in society. The features of this interval that applies to systemic risk, in addition, should have very specific properties. What we can call the acceptable systemic risk interval should be so constructed that it affects option risk in a specific way. The systemic risk interval should be such that all exercises of option risk, even for those who are ‘the least well off’ in terms of their ability to manage systemic risk, can be plausibly depicted as autonomous decisions that are compatible with a person’s self-respect. As we have stated above, we take no specific stance about where exactly to situate the systemic risk interval. Nonetheless, we will say something more about the reasoning that explains its identification.

While we cannot purport to offer a precise characterization of the features of what we have called the acceptable systemic risk interval, we can nonetheless provide ‘markers’ that would enable one to characterize such an interval. Following Baderin and Barnes (2018), we suggest that acceptable systemic risk as we have depicted it above describes an option set in which decisions about which risks to take and how to hedge against risk are made in epistemically favorable circumstances, where risks inevitably born (i.e. those that are not voluntarily undertaken) are non-pervasive, and lastly, where the stakes involved in such risks are not uniformly high. The systemic risk interval, as any interval, has two boundaries. The lower boundary, as we have seen, is characterized by the point at which less opportunity to exercise option risk entails what we have, so far quite vaguely, defined as a dull life. In what follows we will try to show that when a person’s option set is too constrained, that is, when systemic risk is too low and does not really allow her to exercise meaningful option risk, her self-respect might be endangered.

Let us then consider this lower bound on systemic risk. While we have offered examples that of risks that enrich one’s life, these are often option risks. Leaving one’s house, going after a particular career, or asking someone on a date are all risks, but those are risks borne by individual choices. Why might some systemic risk be good?

One clear reason is that when systemic risk becomes too low, the very possibility to exercise our choices vis a vis option risk are curtailed. And, no possibility to exercise option risk, at the limit, translates into a life in which persons are not agents - one in which they do not make enough of a contribution to their lives going well. Adapting Gheaus and Herzog’s (2016) view of the ‘goods of work’, we can say that the opportunity to test one’s ability to manage risk is also important to one’s self-esteem, and thus, to one’s self-respect. Mastery of risk management can empower individuals by giving them the sense that ‘they are able to deal with what life throws at them’. The ability to manage risk can only be acquired by being regularly exposed

to it, and it is only by acquiring the latter skill that we can lessen our fear of uncertainty. As there is no hiding from the precariousness of the human condition, its acceptance and the development of the skills that can allow us to navigate it are a central good for any individual.

A core reason for valuing systemic risk exposure is thus that it promotes the development of agency. We can find a parallel argument to that of Mill's in chapter 4 of *On Liberty* (2011\1859). As Mill notes, even if it were true that paternalists knew what was best for an individual, the individual still suffers a loss if she never develops a capacity to judge for herself what is best. Systemic risks expose us to choices with real consequences. This is important both for the consequence and for the fact of choice. If we mitigate or eliminate risks too much, then people can't develop their capacity for judgment, nor can they develop a sense of self-respect that comes from guiding their own life. Imagine a country in which everyone receives the exact same income, for instance. Though this protects people from the risk of an unsuccessful career (or more seriously, homelessness), it also denies people a reason to develop their capacities.¹¹ While many may well be driven to achieve their goals regardless of reward, plenty will choose leisure over work or education. Ignore the productive costs, and instead focus on the individual loss. Those individuals are harmed by never needing to make substantive choices that shape who they become. While we need not go so far as to say that exposure to *devastating* risk is needed, we do wish to say that if significant outcomes in one's life are independent of one's choices, then individuals are not leading lives that are their own. Protecting people too much denies their ability to express their values and goals through their own choice-making. Note, in addition, that in our view this is the case even when the aforementioned 'protection' would promote the long-term interests of individuals. This is so because being able to make 'mistakes' is often a crucial element of one's identity. To be able to select specific areas in which one's preferences might set back one's interests contributes to making us the persons we are. To illustrate, someone who has immense passion for food might decide that being overweight, while bad from a health perspective, allows them to more thoroughly explore his lifelong passion for sophisticated cuisine. An all-knowing, benevolent and perfectly efficient public health system might protect the individual in question from ill health in a comprehensive way, but in so doing, it would also fundamentally alter who they are.

The upper boundary of the interval, as we have seen can be depicted as a situation in which individuals are forced to exercise option risk all the time, and thus

¹¹ Some would deny this claim. Much will depend on the level of idealization one is ready to accept when doing normative theorizing. A relatively deflationary answer is that, for those that are prepared to accept a greater amount of idealization than we are, then, they should understand what we are doing as discussing a normative theory for a non-ideal world.

important aspects of their lives are constantly up for grab. A life in which we have no option but to undertake risky choices, and such choices are relevant when it comes to our most valued ends, is not, we will argue, an autonomous one. This is for several reasons. When everything is constantly up for grabs, we have less time to think, and we are more concerned that every possible decision we make may have devastating consequences for our future. When this is the case, the quality of our decision-making will tend to degrade, and thus the choices we make will not be free and informed. This is bound to affect our view of their quality – epistemically bad choices can be autonomous only up to a point. Second, following Raz (1986), autonomous choice presupposes a menu of valuable options to select from. And this menu, we believe, must contain some option not to bear risk at least some of the time, and at least concerning some central elements of one's life.

Just as going below the lower bound of the systemic risk interval would lead to an unsatisfying life, going above the upper bound of the systemic risk interval would lead to a life that could hardly be characterized as a manifestation of one's own agency. Choices are difficult to evaluate if the background conditions are in constant flux. As a simple example, one can't choose to take up a night class to better their condition, or even arrange for childcare if their employer doesn't set work schedules in advance. Working with an "on demand" schedule makes planning impossible. Many service industry jobs in the United States work this way, now that more efficient scheduling systems exist. But this "just in time" approach to scheduling introduces significant risks on a segment of the population that is less well equipped to absorb extra risks. Indeed, it makes the exercise of agency in other areas of one's life far more difficult.

Three features of our analysis are worth mentioning. These features concern the broader context of our view. First, we would like to stress that our view is weakly comprehensive. It is comprehensive because it endorses, much like classical liberalism, the value of individual autonomy. It is weakly so because autonomy is not, in our view, one master value that comes to dominate other, more political, values. Instead we understand autonomy as valuable for two reasons. For example, it normatively grounds individual choice. Without some measure of autonomy, choice lacks normative content. How much systemic risk a society imposes on its less well-off individuals and, more broadly, the kinds of option sets faced by persons, affects the kinds of citizens they will be. Active political participation, and a secure sense of one's worth, requires citizens to be placed in an environment in which they are neither completely sheltered from risk nor pervasively exposed to it. In addition, our attention in this paper is focused on levels of systemic risk. However, we do not wish to claim that systemic risk is all that should matter to a theory of distributive justice. How societies deal with systemic risk is but one element of a complete theory of justice. Finally, note that, as we have argued time and again, at least some of the concerns with levels of systemic risk are connected to the kind of society that we

can collectively hope to aspire to in terms of levels of innovation and economic growth. While innovation and growth are not necessarily intrinsically valuable, there is a clear sense in which we can give a non-comprehensive justification for their public significance. More innovation means, *inter alia*, a greater ability to cure illnesses, and help people that suffer from chronic medical conditions. More economic growth may (but of course need not), for example, allow greater investments in human capital leading to better jobs for more people, and to a more aware citizenry (see Acemoglu, 2019).

Third, we believe that our analysis highlights some of the shortcomings of both the high liberal and classical liberal tradition. Put differently, considering the amount of systemic risk present in a social system offers a new way at looking at accounts of distributive justice offered by both the high liberal tradition and the classical liberal tradition. For both high liberals and classical liberals, there is a need to take into account the kind of social system that their conceptions of distributive justice allow us to imagine in light of the amount of risky choices that such systems would require individuals to make.

Stated more precisely using the language we have developed in the paper, we can think of these two traditions as only focusing on one portion of the systemic risk interval. High liberals are concerned by the effects of excessive systemic risk on individuals but tend to forget that too little of it and the option set with which persons are confronted might not allow them to deploy their agential capacities in normatively relevant ways. Classical liberals tend to be sensitive to the latter problem but then seem to forget that there must be restrictions on the option set faced by individuals if their choices are to be depicted as truly autonomous. More systemic risk means more opportunities for choice, but not all such opportunities lead to more autonomy; if systemic risk swamps one's ability to take on option risk, then, autonomy is inevitably curtailed. In addition, both high liberals and classical liberals should also take into account how levels of systemic risk are (or are not) compatible with a dynamic society, one that is conducive to technological progress and thus economic growth. While the latter is a largely instrumental consideration, we believe that it is a central one; especially in light of the effects that technological change and growth can have on human welfare, poverty reduction, and social mobility (see Cowen, 2011).

6. CONCLUSION

In this paper we have argued for a view we have labelled 'interval limitarianism' applied to the idea of systemic risk. The view suggests that social systems need to be so structured that the level of systemic risk they ask individuals to bear, especially those that are less well equipped to manage it, should stay within two boundaries (*i.e.* the two limits of the interval). Less of it, and persons will, at the limit, stop being

agents, relinquishing part of their self-respect in the process. Collectively, a society that features too low a level of systemic risk will be stagnant and won't be able to produce sustained innovation and growth. Yet, at the same time, allow a social system to feature excessive amounts of systemic risk and other kinds of ills will appear. Individuals will become less able to pursue their conception of the good over time and, at the limit, will stop leading autonomous lives. Collectively, a society where systemic risk is too high will be one where the ability to invest in the future is stunted by the continual need to manage the present. Thus, systemic risk, we have claimed, is much like medicine: a social system needs just the right amount. The latter is, in our view, an intuitively attractive conclusion to reach. This helps us make sense of the somewhat intractable nature of the dispute between classical and high liberals, as each has focused on one end of the interval. Between the Scylla of a fully planned life where all risks are managed by others on our behalf, and the Charybdis of a world in which uncertainty degrades our ability to choose autonomously lies the hope for some form of reconciliation. The hope of taking risks with our lives knowing that at least some parts of it are not up for grabs.

BIBLIOGRAPHY

Acemoglu, D. (2019). *It's Good Jobs, Stupid*. Economics for Shared Prosperity, Policy Brief 13. Available at <https://econfp.org/policy-brief/its-good-jobs-stupid/#> (accessed on October 23rd, 2019).

Adams, John, (1995). *Risk*, London: Routledge.

Adler, D. and Posner, E. A. (2006). *New Foundations for Cost-benefit Analysis*. Cambridge: Harvard University Press.

Baderin, A. and Barnes, L. (2018). "Risk and Self-Respect". *British Journal of Political Science*, (in press) pages 1–19.

Barr, N. (2012). *Economics of the Welfare State*. Oxford: Oxford University Press.

Cohen, A. and Einav, L. (2007). "Estimating Risk Preferences from Deductible Choice." *American Economic Review* 97(3): 745–788.

Cowen, T. (2011). *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*. London: Penguin.

Cowen, T. (2017). *The Complacent Class: The Self-Defeating Quest For the American Dream*, St. Martin's Press.

Freeman, S., (2011). "Capitalism in the Classical and High Liberal Traditions", *Social Philosophy & Policy*, 28(2): 19–55.

Frick, Johann. "Contractualism and Social Risk". *Philosophy & Public Affairs* 43.3 (2015): 175–223.

Gheaus, A. and Herzog, L., (2016). "The Goods of Work (Other Than Money!)", *Journal of Social Philosophy*, 47 (1): 70–89.

- Hansson, S. O. (2013). *The Ethics of Risk: Ethical analysis in an uncertain world*, New York: Palgrave MacMillan.
- Hansson, S. O. (2018). *Risk*. Entry for the Stanford Encyclopedia of Philosophy. Zalta, E.N. (ed.). Available at <https://plato.stanford.edu/entries/risk/>. Accessed on Wednesday October 10th 2019.
- Hacker, J. S. (2006). *The Great Risk Shift: The New Economic Insecurity and the Decline of the American Dream*. Oxford: Oxford University Press.
- Heath, J. (2001). *The Efficient Society: Why Canada is as Close to Utopia as it Get*. New York: Penguin Random House.
- Hermansson, H. and Hansson, S.O. (2007). "A Three-Party Model Tool for Ethical Risk Analysis", *Risk Management*, 9(3): 129-144.
- Landes, X. (2015). "How Fair is Actuarial Fairness?", *Journal of Business Ethics*, 128(3): 519-533.
- Lehman, J., (2008). "Contractualism and Risk Imposition", *Politics, Philosophy and Economics*, 7(1): 99-122.
- Mill, J.S. (2011\1859). *On Liberty*. Cambridge: Cambridge University Press.
- Mullainathan, S. and Shafir, E. (2013). *Scarcity: The New Science of Having Less and How it Defines Our Lives*. London: Picador (McMillan).
- Nozick, R. (1974). *Anarchy, State and Utopia*. New York: Basic Books.
- O'Neill, M. and Williamson, T. (eds.) (2012). *Property-Owning Democracy: Rawls and Beyond*. London: Wiley-Blackwell.
- Parfit, D. (2002). "Another Defense of the Priority View". *Utilitas*, 24(03): 399-440.
- Rawls, J. (1999). *A Theory of Justice* (2nd edition). Cambridge: Harvard University Press.
- Rawls, J. (2001). *Justice as Fairness: A Restatement*. Cambridge: Harvard University Press.
- Raz, J. (1986). *The Morality of Freedom*, Oxford: Oxford University Press.
- Robeyns, I. (2017a). 'Having too much', in: J. Knight and M. Schwarzberg (eds), *Wealth: NOMOS LVI*, NYU Press, pp. 1-44.
- Robeyns, I. (2017a). *Wellbeing, Freedom and Social Justice: The Capability Approach Re-examined*. Cambridge: Open Book Publishers.
- Schumpeter, J.A. (1994). *Capitalism, Socialism and Democracy* (5th edition). London: Routledge.
- Sen, A. (1999). *Development as Freedom*. New York: Alfred Knopf.
- Sen, A. (2009) *The Idea of Justice*. New York: Belknap Press.
- Tomasi, J. (2012). *Free Market Fairness*. Princeton: Princeton University Press.