The Market and the Dystopian Tradeoff

Nicola Gennaioli[1]

April 10th 2020

Weeks into the COVID crisis, the economic costs of lockdown are clear. Production collapses in many sectors, incomes are lost, even subsistence is at stake for many. Contagion and hospital congestion have gone down, but the virus is still around, potentially for months. In this doom state of affairs, many observers propose ways to restart our economies. The main approach focuses on health policy, particularly on designing testing strategies to isolate the infected and reduce contagion. Testing is a first order component of any strategy, but probably is not enough. At a minimum, it will take considerable time to scale our testing capacity to the level required for the virus to disappear. A complementary approach, then, argues that cohabitating with the virus means facing the so called "dystopian tradeoff": the choice between economic growth and lives lost. This is part of the calculus, but it does not say how we optimize the economy even conditional on the number of deaths we may be ready to accept.

To make progress, it may be useful to take a step back and think about how an economy infected by the virus works. Such an economy is in many ways special, also in the challenges it faces. Sure, a "viral economy" is infested by two negative externalities: contagion and hospital congestion. To soften them, the government should reduce the level of activity. But how do these externalities depend on individual decisions to trade and market incentives? How do individual decisions adapt to these externalities? I offer some preliminary thoughts on these questions. I focus on an important yet in my view insufficiently recognized aspect: heterogeneity in externalities. By this, I mean heterogeneity in the contagion risk posed by different technologies, and heterogeneity in the contagion risk tolerance of different people. Such heterogeneity, I shall argue, is key to think about how the viral economy works, how severe the overall externalities are, and creates difficult tradeoffs. Taking this into account may help practical thinking.

Consider heterogeneity across activities first. Massive gatherings at sport events, concerts, or the concentration of kids in schools or universities are conducive to much more transmission than other social activities, such as small parties among friends. Food and drugs are essential for survival but large parts of these industries' supply chains take place in establishments that pose significant contagion risk. Services such as bars, restaurants, hair salons, etc., may create even larger risks. Critically, in virtually all sectors it is possible to employ safer technologies. Most educational activities can be carried out online, which allows otherwise risky activities to pose near zero contagion risk. Some safe technologies are cheaper than traditional ones because they use less commercial space, but the quality is lower (e.g. home delivery of grocery and food). Safer technologies are often more expensive: they rely on greater distancing, use of thermometers, etc. We could of course see significant creativity in the invention of new, economically more efficient, safe techniques. How do we choose which technologies to operate and to what extent? How do we tradeoff safety, costs and quality? How do we support innovation?

Second, individuals vastly differ both in their tolerance of contagion risk and in the contagion risk they pose for others. Based on current information COVID appears to be much less lethal for the young and perhaps less lethal for women than for the average person, and more lethal for the elderly or the already sick. As we gather more data, we will surely learn more about variation in mortality risk. But there is strong heterogeneity in the willingness to bear contagion risk, too. On the one hand, individuals vastly differ in their risk tolerance. Some people avoid taking airplanes or buy only safe assets, others invest in risky assets

[1] Bocconi University and IGIER. I thank Daniele D'Arienzo, Daniele Goffi, Nicola Limodio, Alberto Martin, Giacomo Ponzetto, Stefano Rossi, Andrei Shleifer, Guido Tabellini, Marco Tabellini and Jaume Ventura for helpful comments.

and work in dangerous occupations. On the other hand, the cost of dying also differs among people, for instance due to differences in productive capacity and happiness. People who tolerate contagion risk arguably take fewer precautions, so they also create larger negative externalities for others. But there are also large individual differences in the very cost of taking precautions. Some people are fond of social interactions or dislike washing their hands, others barely tolerate social contact and are mysophobic. The latter people naturally take more precautions, so they are safer to interact with than the former. How do we make sure that different people take privately and socially optimal decisions? How do we distribute jobs and externalities across them? Cost-benefit analyses that neglect this vast heterogeneity are problematic.

Based on this premise, I argue that optimizing a viral economy does not just entail reducing the level of activity. It also means implementing an optimal: i) mix of safe and risky technologies, and ii) allocation of contagion risk across people. The optimal allocation depends, of course, on the available technologies and on the preferences/mortality risk of individuals. To the largest possible extent, contagion risk should be shed on the best bearers of it, but the overall prevalence of contagion depends on the availability and cost of safe technologies as well as, crucially, on testing-based health policies. Recognizing this enriches but in some sense challenges the logic of the "dystopian tradeoff": an economically efficient allocation can have great health benefits. The average mortality rate could be orders of magnitude lower if contagion risk is concentrated on low mortality risk people. This means that promoting a more efficient allocation of risk could also allow to reduce the overall level of restrictions, reducing economic costs.

Of course, reaching an efficient allocation is difficult. It is based on dispersed knowledge about preferences and technologies. It is also based on dispersed medical expertise and creativity that, combined with entrepreneurial skills may promote safer and more efficient methods of production and trade. Which allocation mechanism should society employ?

I first consider how a competitive viral market may work. I argue that such market faces two challenges. First, some risk tolerant people take socially excessive risks. Second, some risk intolerant people distrust social interactions and retract from trade. Risk tolerant people exacerbate contagion, which causes deaths but also distrust, creating economic costs. I argue that even if the market may produce a socially excessive level activity, it softens these problems by creating sorting: risk intolerant types tend to sort into safer activities, risk tolerant types into riskier ones. In this allocation, risk intolerant types accept lower real wages (they work/consume using costlier/less enjoyable technologies) but suffer from fewer negative externalities than risk tolerant types. By contrast, risk tolerant types engage in more lucrative/enjoyable activities, but also take higher contagion risk, which they are better equipped to bear. Due to sorting, the distribution of externalities is uneven.

The market induces sorting by pinning down "the price of safety". This is captured by the wage discount in safer jobs and by the price premium of safer goods. It captures both the availability of safe technologies as well as the extent of society's tolerance for contagion risk.

I then argue that the market under-provides sorting. In part, sorting is limited due to technological constraints (e.g. contacts in public sidewalks). A deeper reason is that sorting is limited due to individual incentives to mix across safe and risky activities. Mixing is problematic in two ways: it sheds too much risk onto the risk intolerant types who trade, and creates strong distrust in other risk intolerant types, who refrain from trade. In fact, these types may end up trading too little.

Due to mixing, the equilibrium market price of safety is distorted and fails to trigger efficient decisions. Invention and adoption of safer technologies that entail some social contact is not sufficiently rewarded. Self-isolation and technologies avoiding any contact are too rewarded.

The market allocation may be improved by creating mechanisms that facilitate sorting. Following a longstanding tradition in law and economics, this is akin to finding a way to strengthen the enforcement of

property rights on the safety purchased in the market. I discuss a highly stylized example in which voluntary membership in risk-rated groups with minimum safety restrictions may help achieve this. Even if group boundaries are imperfectly enforced, this approach may allow: i) market prices to become more informative, ii) firms and consumers to make more efficient decisions, and iii) the epidemiological dynamics to be more benign, with contagion risk being more concentrated on more efficient bearers of it.

The writing goes as follows. I first discuss how the presence of a virus may affect the market, and the entailed market failures. I then consider how voluntary membership in "safety-rated" groups may help improve matters. I discuss many practical challenges such an organization would face, including individuals' distorted perceptions of their own mortality risks. I conclude by discussing existing approaches to restart the economy in light of my analysis.

A Market Infected by the Virus

The virus injects contagion risk in all human interactions, including market ones. Physical proximity or sharing of objects/surfaces may transmit the disease from person to person. Contagion risk is present during consumption and leisure (e.g., when attending crowded shops or disco clubs), with coworkers, among friends, within families. It is present, and highly lethal, in hospitals. The sender of the virus may be aware or unaware of his infection, the receiver is often unaware of the sender's. Therefore, transmission is almost always involuntary, creating an externality. This externality is quickly invoked to justify extensive government intervention. Before discussing this, consider how it affects an unregulated market.

As the virus becomes salient, people become aware both of contagion risk and of its dangers and change their behavior, even absent any government policy. They start distancing, they purchase and wear masks (if available), they stop going to the restaurant, etc. This response will of course not be uniform. The young, who are fonder of social interactions and less risk averse, will change their behavior less than the elderly, who are also more wary of health risks. Some highly risk tolerant people will continue to engage in very risky activities. Some highly risk intolerant people will self-isolate. To an observer unaware of the virus, this economy would appear unusual. People take many precautions. Not the socially optimal ones: the level of activity will be too high due to externalities. But surely many more precautions than in normal times.

Heterogeneity in risk tolerance does not simply mediate individual responses to contagion risk. It sheds new light on the contagion externality too: it is conducive to an incentive to sort that affects how contagion is spread. In an unregulated market, not all disco clubs would shut down. Those run by risk tolerant managers and attended by risk tolerant patrons will stay open. Conversely, not all groceries will stay open. To avoid the risk tolerant dancers, the highly risk intolerant inhabitants of a neighborhood would only shop online. Local groceries would have no business. Risk tolerant and risk intolerant types would tend to separate and avoid each other.

Sorting would be health efficient. By interacting among themselves, the more contagious risk tolerant people would mostly infect their likes. This is a small externality. Likewise, by interacting among themselves, the little contagious risk intolerant people would be a weak vehicle for the virus. This is also a small externality. Sorting has further benefits. Some risk intolerant people, while restricting interaction, will not fully isolate. They are willing to have some interaction with risk intolerant types they know. Their likes avoid high risk situations, so they are safe to interact with. Such interaction among risk intolerant types will be a bit too high relative to what is optimal under sorting. This is the typical excess level of activity occurring under negative externalities. But it would arguably be better than full isolation.

This observation yields another implication: a store mostly visited by risk intolerant consumers may be quite safe even if it uses a fairly standard technology. Risk intolerant people take more precautions such

as washing hands, wearing mask, gloves, etc. As a result, they are less likely carriers of the virus. Maybe the store does not need to shut down or to buy an expensive "self-isolation" technology. With some precautions, and with sorting of risk intolerant people, it could stay afloat. Sorting among the risk intolerants reduces contagion risk, so it is a source of mutual trust and trade.

Of course, as already hinted, these voluntary mechanisms do not eliminate externalities. But the mere recognition of contagion by consumers will change their behavior in some desirable ways. Overall transmission would go down because of voluntary precautions, contagion risk would be more concentrated on risk tolerant types because of sorting, consumption would shift toward safer technologies such as online grocery and food delivery, and the overall volume of deaths and trade would go down.

Consumers will not be the only ones changing their behavior. Firms will do so, too. Before public lockdowns, some large events such as concerts were voluntarily canceled. Several shops and restaurants voluntarily implemented distancing measures or outright closures. Universities voluntarily switched to a safer online model. These private decisions reflect economic incentives. Events with strong interactions create strong externalities. People, afraid of contagion, naturally avoid them. It is natural for organizers to cancel or postpone them, also for fear of reputational loss. Firms in which workers and customers can interact online, such as Universities, face reasonably low costs of shutdown. It is natural for them to voluntarily close so as to cater to their stakeholders' demand for protection. In so doing, firms spontaneously switch to safer technologies, whose demand has soared.

But market incentives may do even more. As the virus spreads, safety becomes a first order driver of decisions. It becomes a, if not the, key product characteristic. As a result, stores, movie theaters but also manufacturing firms openly compete by implementing yet more rigorous safety: wider distancing, repeated testing of their employees/regular customers, etc. These investments increase the appeal of these facilities for customers and workers. Competition in safety may spur valuable innovation. We learned that there is a lot of dispersed knowledge and creativity in society about how to tackle contagion risk. Many countries with federal health systems experienced vast regional differences in the efficacy of public health policies. Some regions – such as Veneto in Italy – relied on local experts to develop innovative and highly effective mitigation policies, teaching key lessons not only to their peers, but to the central government itself. Large firms also invest and innovate in safety, by consulting with medical experts and epidemiologists. This is another important margin of the market's response.

In the unregulated market, this myriad of individual decisions about which technologies to adopt and which jobs/stores to choose will eventually affect prices. Risk intolerant workers will accept lower wages to operate costlier but safer technologies. Likewise, consumers will accept higher prices to buy safer goods and services, stimulating their supply. The ensuing wage differential between risky and safe jobs and the related price differential between safe and risky goods have an important economic meaning. They capture the market's price of safety. They offer an indication of the extent to which society is willing to pay to reduce the possibility of contagion in labor and goods markets. They guide the decisions of many consumers on whether or not to take some risks when shopping, the decisions of many workers on whether to take a riskier job, and it pins down the relative profitability of safer vs riskier technologies, influencing entrepreneurs' decisions of which ones to adopt or innovate.

For any given price of safety, relatively more risk tolerant workers will sort into risky technologies, more risk intolerant ones sort into safer technologies. If safe technologies are too expensive, wages in safe activities will be relatively low and prices of safe consumption goods will be relatively high. Buying safety is expensive. Not a nice situation during an epidemic, but one that can be hardly avoided without technological innovation. Coming up with such an innovation would be fantastically rewarding, given society's willingness to pay for it. More optimistically, as the virus recedes, demand for cheaper riskier goods will increase, reducing the price premium carried by safe goods. Likewise, labor supply to risky jobs

will increase, reducing the wage premium they carry. The reduction in contagion risk reduces the market price of safety, signaling that society is ready to go back to "normal".

The Contagion Externality

Do these arguments imply that unregulated competition will trigger an efficient allocation? The answer is of course not due to externalities. One source of externalities is that private measures are partial. Common spaces such as sidewalks or squares are loci of contagion and falls outside the remits of private decisions. Although concerning, in these locations contagion risk is milder because the encounters occur outdoor, are quick, for the most part do not involve contact, and they can be anyhow regulated by public authorities. The riskiest encounters occur indoor, where ownership is more easily demarcated.

The chief reasons why the market fails are two. The first is incentives: as much as we all try to avoid the virus, we are much less concerned about the risk of spreading it onto someone else. This implies that, whichever society or subgroup we are thinking about, the possibility of contagion tends to induce a socially excessive *level* of market activity (and insufficient precautions). The second, and key reason is lack of knowledge: we do not know who is a contagion risk taker and who is a contagion risk avoider. This implies that not only activity will be inefficiently high, but the externality may be very severe too, because risk tolerant types can infect risk intolerant ones. If contagion was especially prevalent among high risk tolerant people, externalities would be small despite a high activity level. I call "mixing" the presence of different risk types in the same facility. As we have just seen, mixing is a key reason for market failure.

Mixing is perversely intertwined with market incentives. We may choose to buy basic goods in a cheap store that implements no safety measures, but subsequently purchase differentiated goods in an expensive store that implements safety restrictions. In the first store, we incur contagion risk for personal monetary benefit. In the second store, we shed the risk, unwanted, onto safety-seeking consumers. Yes, perhaps the very same differentiated goods may be sold at discount by another, low safety, store trying to attract risk tolerant customers. But one cannot expect full sorting to occur for at least two reasons.

First, many aspects beyond safety, price and quality may render a store an imperfect substitute for another. A risk intolerant type may shop in a cheaper store because of its physical proximity. When meeting risk intolerant shoppers in a safe store, he could infect them.

Second, somewhat risk tolerant consumers may wish to partially insure by optimally dividing their shopping between risky and safe places, again generating large negative externalities in the latter. Behavior in labor and goods markets may have the same effect: people could reap low prices by consuming in risky stores, while accepting low wages from working in safe jobs (or vice versa). This mismatch between production and consumption activities messes up with sorting, creating large externalities.

As I already argued, mixing misallocates contagion risk and creates distrust by risk intolerant types. But from a market viewpoint, mixing renders prices less "informative" of the social demand for safety. The wage cut I accept for working with a safe technology cannot be so large if my coworkers are risk takers. The technology may protect me during production, but significant contact with coworkers remains during non-production activities in the firm's premises. Likewise, the price premium I am willing to pay for consuming at a safe store cannot be so large if the store is also visited by risk takers. Subtly, because mixing reduces price and wage gaps, it induces even more mixing. As the market price of safety falls, more risk tolerant types will be willing to engage in safer activities, causing the price to fall even further.

This mechanism distorts technology choice and investment. There would too small a price premium for safer technologies that still rely on some social interaction. Their adoption/invention is stifled. The price

premium would instead be too large for activities that do away with social interaction altogether. Their adoption/invention is excessive. Because social interaction is valuable, this may be a problem.

The contagion externality therefore inhibits the working of the market. Sorting is insufficient, health risk is inefficiently allocated. Trust in trade is low, and market incentives to invest across different safe technologies may be distorted. The level of production may be too high, due to non-internalization of contagion, or too low, due to the mistrust that contagion creates on the risk intolerants.

The Congestion Externality

Another key externality during an epidemic works through hospital congestion. Widespread contagion may cause many patients to simultaneously need intensive care. This increases mortality risk because of unavailability of treatment for all. It may also enhance viral transmission in hospitals, shedding contagion risk on highly risk intolerant people. The market of course does not tackle this externality, because no "hospital congestion" price is attached to individual risk taking. This is even more evident after we recognize that societies, even those with private health care, do not accept to let the sick die unattended. The market may soften congestion through sorting, because less contagion on risk intolerant people may free hospital beds. But it may not do so, for it could increase the number of risk tolerant types needing health care. To tackle congestion effectively, the level of activity must also be targeted.

Other Market Failures

It could be argued that during an epidemic there are other, potentially even more severe, market failures. Some stress that the market is unable to quickly reorganize productive units or collect financial resources to produce medical devices or safe technologies. There may be some truth to these concerns, but in part these failures may be due to the market's inability to properly price safety, in part because of mixing, in part because of other price frictions. Other market failures may be wholly different in nature and unrelated to the epidemic. They may have to do with general contracting problems and misaligned incentives of privately vs publicly owned firms. While potentially important, these issues are separate from my main focus here, so I abstract from them.

Enabling Sorting

I stressed that negative externalities during a pandemic are aggravated by excessive mixing of risk tolerant and risk intolerant types into the same activities. It may therefore be helpful to facilitate sorting based on risk tolerance. The importance of sorting is self-evident. At a Champions League game, many young spectators from Bergamo got infected and contaminated fragile elders at home. This is a drastic manifestations of mixing. Disastrous mixing took place in hospitals and hospices, where infected patients transmitted their disease to patients with other diseases. One lesson is that hospitals and similar facilities should be clearly separated from society, and their personnel should carry guarded lifestyles.

Can we create a mechanism enabling more, even if partial, sorting also in the market? In the abstract market I just analyzed, greater sorting could be harnessed by allowing voluntary membership into different safety ratings. Each safety rating is characterized by a different minimum degree of precautions.

Sorting into Safety Ratings

To see this vividly, imagine a "traffic light" model with three ratings: red, orange and green. Red captures the riskiest activities allowed, such as working in properly distanced and sanitized manufacturing plants or services, consuming in similarly safe shops and bars, and perhaps – if the epidemic is contained – attending highly spaced public events such as football matches. Green captures the absolutely safest activities such as staying at home and working from home, shopping in stores where few customers are allowed, or working at firms where very strong precautions are taken such as aggressive distancing, periodic testing, thermometers, etc. Green could be similar to current lockdowns, even though activities that are currently banned could be opened provided they adopt a "green" technology. Orange captures middle of the road safety, such as theatres keeping a certain minimum spacing among seats, and firms employing middle of the road safety. Public parks could be divided into red, orange, and green depending on the activities allowed. Public transports could also be divided into red orange and green areas.

Think now of attaching a risk rating to economic agents, too: red, orange, and green. Participation into a group is voluntary. A person can only interact with people of his own color and engage in activities of the same color. There could be inter-group trade, but only for interactions occurring at a distance, such as when goods and services can be safely shipped without creating contagion. Membership should be visible (e.g. wearing a bracelet that emits colored light), to ensure that matches occur at least along mutually desired lines. There would also be a fourth, highly precious, white group. These are the people who have contracted the disease, have recovered, and now are (temporarily?) immune. They can engage in any activity they want, including normal times ones, and interact with whomever they want.

Entrepreneurs can of course choose which group to join, either by choosing which activity to engage in, or by choosing which technology to operate. Larger firms could organize some red shifts/plants, some orange shifts/plants, and they may also have green activities. Within a group, firms can innovate, but only towards safety, not toward risk.

The precise nature of each group will depend on the equilibrium, which in turn depends on the distribution of risk types, their consumption preferences, the available technologies, etc. A plausible outcome is one in which joining the green group means buying safety at the cost of a lower real wage, in particular yielding fewer non-tradeables (green technologies are expensive). Joining the reds exposes you to the largest risk of contagion, but yields a higher real wage and in particular many more non-tradeables, including more social interaction. Finally, becoming orange would mean obtaining middle of the road safety and real wage/non-tradeables.

Trade among groups could be large. Many tradeables might be produced by the red using cheaper technologies. These goods are be shipped to the older, sick, or wimpy greens, who have assets to pay for them as well as income earned by working with green technologies (or obtained from the government). Reds receive higher wages from producing tradeables, which could be exchanged for tradeables produced by the other groups. The orange could produce in safer manufacturing plants. The green could produce online educational services, medicine, or work as plumbers or electricians with appropriate precautions and distancing. Isolated jobs could be performed by the orange or by the red, but the latter would shop in riskier stores, benefitting of course from their lower prices and greater interaction.

In this world, the real wage gap across riskier and safer groups would capture the market price of safety. It would not be a perfect price, it would be too rough/categorical, but it would be more precise than one in which groups are not demarcated. If the real wage premium of riskier groups becomes very high, for instance because there is a shortage of red labor, people would move from the safer groups to the red one. This causes the red wage to drop until the marginal worker is indifferent. If the red premium is unduly low, labor flows the other way around until the red wage is high enough that equilibrium is restored.

Changes in membership are allowed, but with restrictions. Switches from safe to risky groups pose no problem and could be immediate. Switches from risky to safe groups are instead problematic because they can raise contagion among risk intolerants. So these switches are conditioned upon quarantine or a medical test. Capital would be allowed to chase return differentials across groups, but more easily than labor because it carries no infection.

Decisions will not be easy: being a red, orange or green type would affect both work and consumption opportunities. It would be akin to moving to a different region or country. But being in a safer group would allow risk intolerant people to have the desired level of protection and yet to have large enough trust to eschew full isolation, without having to force risk tolerant types into draconian restrictions. By setting minimum safety standards in each group, the government would reduce total contagion risk but it would also affect its distribution. By sorting, safety among the risk intolerant types increases, allowing to marginally relax restrictions to red types relative to what is necessary under a uniform policy.

Ratings and the Contagion Externality

In our funny society, ratings soften the contagion externality by reducing mixing. Because group membership is visible and verifiable, it enables sorting in all structured transactions. Members of safe groups wish to minimize contagion risk (or do not want, in public settings, lose their reputation), so they avoid interactions with risk tolerant types, who carry the visible label of their membership. Stores and firms would install devices to allow entry only by consumers/workers of the appropriate color.

Sorting would influence epidemiological spread. In terms of the conventional SIR model, the transmission parameter R0 would be group-specific. It will be lowest among the greens, it would be higher for the other groups. If the red take little precautions, R0 might be above one for them, creating escalating contagion. This may be undesirable, even if this group is risk tolerant. But crucially, the group structure would help mitigation, for the government could selective test and isolate red people. This may be useful, especially if this society has a low testing capacity.

It may seem funny that sorting could move epidemiological dynamics in a group specific direction. But this already happens in society. Several people argued that in Italy COVID may have especially spread among the risk intolerant elderly because of the integrated family structure. In other countries, such as Germany, the virus especially spread among risk tolerant younger types. These cross country differences suggest that contagion and the severity of the epidemic may be strongly influenced also by the social network in which transactions take place. In my abstract society, the ratings system would temporarily influence the social network in a way that enhances sorting along risk tolerance classes.

With a meaningful spatial dimension, the equilibrium ratings pattern would probably form geographical clusters. These could be due to social imitation, peer pressure, local culture, or technology. Touristic locations on the coast, in islands or the countryside could become “green areas” with strong precautions to which self-sufficient risk intolerant people temporarily move and to which green workers move, too. Locations with better health sectors, or cities, may be predominantly red and orange because of the lower risks, higher density, and greater benefit from social interactions. Locations with worse health sectors or towns in the countryside might be predominantly green due to the greater risks, lower density, and different age structure. Areas that have already experienced a strong epidemic will likely become red too, as they are already populated by many whites. Accordingly, areas that have witnessed highly limited viral spread could become predominantly red, becoming productive centers.

Of course, this system would not fully avoid mixing. One enforcement problem arises when a safe group member has a private incentive to illegally trade or interact with a risky group member. The “safety

traitor" would personally benefit without internalizing the negative externality on his peers. Reassuringly, the very low R0 of the greens would cause this occasional infection to quickly die out. The government could also come up with creative proposals to better enforce group boundaries, such as using gps to track dangerous contact among different colored dots. This would raise privacy concerns, but its use might be made voluntary for the really risk intolerant people (e.g., they could join a dark green group). Another enforcement problem arises due to involuntary interaction in sidewalks or apartment buildings. To avoid the highly risk intolerant to have bad occasional encounters, mayors of densely populated cities could slice activities along the time dimension, creating special times during which only green people can be outside. None of these mechanisms will be perfect because of random encounters, limited monitoring, etc. But perfection is not necessary for this arrangement to help. Even partial protection of property rights can improve sorting, the risk allocation, price formation, and trust.

Ratings and Congestion

Once the rating system is in place, the congestion externality could be dealt with using standard Pigouvian taxation. For instance, members of the red and orange groups may be required to pay a periodic fee (no fee for green people). Such fee would probably but not necessarily higher for the reds, because they take more risk, but it should also depend on individual risk factors. For instance, older and/or sick red or orange people would definitely have to pay a higher fee because they are more likely to require intensive care, so they impose a larger congestion externality. Workers in certain higher contagion risk services may also have to pay a higher fee. The fee could be time varying. Periods during which the virus spreads too much in a group, say the red, the red fee could be raised so as to induce switching to safer groups. Subject, of course, to the rules for changing group membership discussed above.

Some Thorny Issues

Even the quirky citizens of the highly stylized and abstract world I described could raise serious objections to the use of the rating system. One philosophical objection is that society should not allow individuals to trade safety because the resulting outcome would be unfair. Poor people with high mortality risk may have to "accept death for bread", joining riskier groups to earn a living. This would indeed occur if green jobs are few for the unskilled or if they pay a below-subsistence wage. I am sympathetic to this argument, but two counterarguments can be made. First, without a ratings system the poor would also have to go to work. Second, nothing precludes the rating system to be augmented with a fairness-promoting policy such as a means and health tested minimum income.

Another objection is that people have inaccurate risk perceptions and hence a system based on voluntary choice cannot work. It is demonstrated that people make certain systematic errors in their risk perceptions. Even if information campaigns make people aware of their place in mortality tables, the young could be excessively optimistic compared to the elderly, and the elderly could be too scared. In fact, excessive distrust in the market could be even due to over-reaction to the virus, an outcome we cannot certainly exclude. Risk assessments may also be unstable over time, oscillating from periods of panic during the cold season, when contagion risk is more salient, to periods of risk neglect in the summer if the virus retreats then. A mechanism factoring in these possible misperceptions would be better. At the same time, notwithstanding errors in perception, a flexible rating system may be better than a uniform policy that entirely rests on the government. Public decisions may in fact be equally influenced by the misperceptions of politicians or by those of pivotal parties/social groups.

A set of practical objections might concern the role of institutions. How can we separate people into different ratings if society already clusters in families, schools, and firms? Take families: given the rules of the game, joining one group or another would be a joint decision. In some cases, all family members could choose to be, say, orange, after discussing the risks and opportunities. In other cases, one family member may leave the others and temporarily join a different risk group to achieve e better family income or safety profile. These decisions will be difficult, but families normally make joint decisions of this sort. The same is true for firms and schools. Educational firms can benefit from the option of online schooling, production firms can separate shifts according to color, or bargain wages with different types of workers. The tradeoffs are undoubtedly complicated. The key point of voluntary ratings is that they move the tradeoffs downstream. Centralized public policies do not remove the tradeoffs. They move them upstream.

Another practical objection concerns market level externalities and fixed costs. What if one group, say the red, turns empty due to uncertainty as to whether there will be any activity there? To deal with this, the government could perhaps engineer subsidies to ensure minimal activity.

It could be objected that it is not desirable to split society into demarcated groups. The politics may be unpredictable (e.g. the red could be stereotyped as "viral spreaders").

I have considered how an equilibrium can be reached starting from scratch. What would happen in a more realistic word that must transition from an old to a new allocation? How to make sure that existing long term contracts are temporarily suspended or rescinded, if needed? This is a huge issue societies are currently facing, for instance with respect to firms' financing arrangements. Currently, this is partly dealt with bargaining and partly it is dealt by having the government (explicitly or implicitly) intervene into existing contracts/rules. This seems natural to expect after a large unexpected shock.

Concluding Remarks

During the quarantine, the COVID epidemic led me to think about how an economy infected by a virus may work. Here I report my limited and preliminary thoughts to the benefit of interested readers. The current emergency may soon disappear or be minimized by the discovery of a cure, in which case we will all resume to normal or almost normal life, and we will forget about this. Or perhaps we will be amused by the madness of these times and of writings like this one. At the same time, why not to consider the worst case scenario of a long journey with the virus. After all, the opportunity cost of thinking is low when you are locked at home, especially if you have a bad internet connection.

Living through the extreme, even if so far brief, Italian quarantine made me appreciate the importance of freedom. I grew up during extraordinary free and quiet times. The epidemic made me realize how important is freedom for life to be meaningful, and how suffocating it is to watch the news waiting for detailed orders about what can and cannot be done. To me, learning about the value of freedom is the silver lining of this unfortunate situation. It motivated me to think about how a free economic system would work when the virus is around, and which challenges it faces.

My wishful conclusion is that even though in a pandemic we must accept restrictions to control the overall level of activities/externalities, we need significantly more freedom than we currently have (at least in Italy) so that externalities are efficiently allocated and a better technology mix is used. In my abstract world, voluntary sorting into risk rated groups with minimum safety levels was the way to achieve this.

I wish to conclude discussing current policy proposals to "restart" in light of my analysis. Several proposals envision selective lockdown for the elderly/return to work for certain low risk people (e.g. the young of women). My analysis suggests that these solutions go in the right direction, but are too blunt, if not brutal. They either force risk intolerant people to self-isolate, or expose them to way too much risk in

daily trades. Neither is desirable. It may be more desirable to create safer circuits in which risk intolerant people (and only them) can interact. This would enhance trust and allow a more meaningful life, without taking much extra risk. But these circuits should be voluntary, for old age is neither necessary nor sufficient for being risk intolerant. By enhancing the safety of vulnerable people, mechanisms like this may allow to relax a bit the restrictions the risk tolerants, to their benefit too.

Another key approach to restart the economy is to do extensive testing. In some schemes, testing is viewed as a way to reduce overall contagion risk, without necessarily affecting its distribution (e.g. random testing). My analysis suggests that these proposals are problematic, because risk allocation is key. Testing everybody means wasting a lot of resources on types who are unlikely carriers of the virus. Other proposals stress targeted testing of people who are central in the contagion network, such as policemen, airport staff, workers at large firms etc. Targeting ex-ante high risk people fits well with my analysis. Finding a way to have ex-ante high risk individuals self-select into testing may fit even better. Testing alone, though, is unlikely to create enough trust until a very large number of tests can be performed.

One broad debate during the crisis concerns the role of the government. In the status quo, the government unilaterally decides what is opened, what is closed, which technology is used, and extensively regulates social interactions. This mode of operation may be justified as a quick top down response to an emerging health crisis. But we may also see heavy public intervention down the line, as we restart our economies. Several observers argue that even in such phase the market will have serious pitfalls, so that the government must directly intervene in allocations, also through public ownership.

One problem with government interventions of this sort is their arbitrariness. In Italy we cannot lift our spirits with a tiramisu because confectioneries are closed, but smokers can because tobacconists are open. This seems funny, but it has strongly different implications for the lives of the confectioner and his staff and that the tobacconist (and also for their customers). This problem may become much more severe as we go along. What kind of caretaking should we allow for the elderly, who have the greatest need but may also be subject to the greatest contagion risk? Is it better to make safety investments in production, in transport, in retail, or all? How do we make sure that safe and cheap technologies are invented? The social value of these innovations may be huge, also in future epidemics. My analysis suggests that the market may induce inappropriate decisions because – due to contagion and mixing – certain safety investments may be mispriced. But it is perhaps better to try to soften contagion and mixing than to do away with the market altogether. Uniform and centralized government decisions will miserably fail to harness highly valuable dispersed information available in the market. Blunt uniform policies may even lead, in the presence of vast individual level heterogeneity, to massive black market activity and corruption.

A more flexible approach for the government would be to use a mix of market instruments. For instance, a not too stringent minimum safety standard could be combined with subsidies to the adoption of safer technologies. This policy seems well intentioned. The safety standard could allow risk tolerant types to have some freedom. At the same time, subsidies to safe activities would enhance protection for risk intolerant people. My analysis indicates that this approach has one key problem: it does not address insufficient sorting. In fact, it makes sorting even more insufficient: subsidies to safe techniques would reduce the market price of safety. As a result, safer jobs/stores would be crowded by many risk tolerant types, generating a very bad externality on the risk intolerant ones. To control this externality, then, the minimum safety standard would have to be restricted.

The bottom line is that, as much as we are all focused on reducing the level of activity, its distribution is equally if not more important for thinking about how to run the economy during a pandemic. To avoid highly inefficient stark and uniform policies, to avoid persistent self-isolation of large parts of society, economic policy should try to reduce mixing, and certainly not increase it.