THE BROOKINGS INSTITUTION
WEBINAR

AI, INNOVATION, AND WELFARE:
A CONVERSATION WITH JOSEPH E. STIGLITZ

Washington, D.C.
Monday, March 7, 2022

PARTICIPANTS:

ANTON KORINEK, Moderator
David M. Rubenstein Fellow, Center on Regulation and Markets
The Brookings Institution

JOSEPH E. STIGLITZ
Professor
Columbia University

* * * * *

P R O C E E D I N G S

MR. KORINEK: Hello, and welcome, everybody. I'm Anton Korinek. I'm a Rubenstein fellow at the Center on Regulation and Markets at Brookings and a professor of economics at the University of Virginia.

I am pleased and honored to welcome all of you to our fireside chat today, and I am especially happy to welcome Joe Stiglitz to our virtual studio. Our topic today is AI, Innovation and Welfare.

Joseph Stiglitz is a university professor at Columbia and he's also the chief economist of the Roosevelt Institute, among many other positions of thought leadership. He is a recipient of the 2001 Nobel Prize in Economics and a former chief economist of the World Bank and chairman of the U.S. Council of Economic Advisors. Joe has made pioneering contributing throughout all subfields of economics. But over the past decade he has been particularly active in research and technological progress and inequality. I have had the honor of collaborating with Joe on several papers in this area and two of them are listed on the web page from which you are viewing this event.

In the interest of full disclosure, Joe was also my dissertation advisor.

Before we start, let me thank the team that has made today's event possible. Sanjay Patnaik, the director of the Brookings Center on Regulation and Markets, and Megan Waring, our events manager.

We have obtained lots of interesting questions from all of you during the sign-up process for this event. Unfortunately, they were too numerous for me to include all of them to our distinguished guest, but I will do my best as we go through the materials to incorporate as many of your questions as possible into our conversation.

The geopolitical context of our event reminds us how much is at stake when we speak about ensuring that technological progress increases welfare, because every time we turn on the news we are reminded that technology can also be used to propagate misinformation and divide us, to suppress people, and even to kill and destroy more

efficiently.

Joe, you have emphasized repeatedly that it's important that we ensure that technological progress benefits everybody in society, not only because that's the ethical thing to aim for, but also because it is crucial for the stability of our society. The stakes are no less than the survival of our western democratic systems and social market economies.

Can you tell us more?

MR. STIGLITZ: A good point place to begin the discussion to realize that over the past 250 years we've had amazing technological progress. Science, the enlightenment institutions have led to higher standards of living, increased longevity. It's hard now to remember, but if you look at the centuries before 1700, 1750, for at least as long as we have data, dating back a couple thousand years, there was no significant increases in standards of living. And yet in the last 250 years it's been enormous. And yet – – and yet, there is not the kind of support that you would have thought for what I sometimes the enlightenment values, the enlightenment institutions, the science, universities, the mechanisms of social, political, and economic organization, rule of law, democracy, separation of powers — a whole set of ideas that we have developed that enable us to cooperate together. Given the benefits, the question is why isn't there more support.

Well, part of the reason there isn't more support is the benefits have not been equally shared. Now, we ought to be recognized, everybody has benefitted. You know, our life expectancy today is so much greater, we just take it for granted. Our standard of living, even the people at the bottom are so much better than they were — than they would have been — if you did a thought experiment — if they were born 250 years. So everybody has benefitted. But some have benefitted so much more than other sand this is contrary to the way the economists typically look at things. We do judge things to some extent in a relative position. So even if everybody has benefitted, if some have benefitted a lot more, the people at the bottom still see themselves as struggling because what they think of as acceptable has moved up with success of our society.

You know, one of the points that has been made by the IMF, by my own research in my book "The Price of Inequality", by a lot of other people, is that when you have a divided society, economic inequality, not only is politics affected, but even the economy may not perform as well. So even from the narrower perspective of economics, not sharing things better —– not perfectly, but better, can have large costs.

And, finally, the point I would make is, well, as an economist, we like to think that everything comes back to economics. We also have to recognize that there are lots of other things going on. That ideas matter, ideologies matter, that if we look at some of the places where populism, nationalism, anti-democratic values, authoritarianism have flourished, include countries that have done relatively well. So you cannot —– there's not a one-to-one like between economic performance and even economic inequalities and the lack of support for these enlightenment values that I just described.

MR. KORINEK: That makes sense. Thank you.

Now, let's start perhaps with the optimistic perspective and let's look at the classic positive effects of progress. The textbook story of technological progress, so to say, and that textbook story is that technological progress benefits everybody, that it lifts all boats. So let's take an example of an innovation in which it seems relatively unambiguous, that it improves people's lives. Say that we invent a new battery that holds more power at less cost and is also better for the environment. Can you walk us through that best-case scenario and through the economics of how an innovation like that would both improve living standards and grow our economy?

MR. STIGLITZ: Yeah, well, let me begin with an even simpler case.

An innovation that produces more output with a given input means that everybody could supply the inputs that they have or work the same amount that they did before and we have more goodies, more goodies to spread among everybody. And so that would be the case where everybody could be made better off. And I say "could be made better off", but that doesn't mean that they will be made better off because whenever you

have a change in technology, it has the potential of changing the balance of market forces. It changes demand and supply curves. That means it changes the market equilibrium, bargaining power. In a competitive market it just changes market prices. That means that it is possible that some groups in the market equilibrium will be actually worse off. So while the pie is bigger, some groups will get so much smaller a slice of that bigger pie that they're actually worse off.

Now, again, in principle, we could take some of the slice, the bigger slice from the gainers and give it to the losers. But there are two problems. One is the guys who've gained may not want to give up and they may have political power. We'll maybe talk about that later. But the economic inequality gets related then to political inequality and that may impede the redistributions that would enable us to ensure that everybody is better off.

There is a second thing that could go wrong, which is redistributions are not costless. We have to have taxes; we have to do other things. And those themselves may make it very difficult to achieve outcomes that make everybody better off.

In addition, I think — I hope we will be able to talk later on, it's not always the case that innovations make society better off. We are drummed — that idea is drummed in our head, but you began the conversation talking a little bit about mis- and disinformation, innovations in military that enhance the ability to kill. There are many other innovations that may increase the profits of somebody, but don't necessarily increase societal well-being.

MR. KORINEK: I want to return to some of those negative effects of innovation again in one moment. But let's stay for a little bit more on the optimistic side. I wanted to ask you, what do you view as the particular benefits that innovations in the area of AI can contribute to our economy? And how can they contribute to social welfare?

MR. STIGLITZ: One of the central aspects of artificial intelligence is that — you can think of this as the third step in the innovation revolutions. The first was that we develop machines that were stronger than us. It could do things physically that we couldn't do. And we would talk about a car having so many horsepower, let alone human power. So

it was physical strength. And then we — in the middle of the last century, we discovered that we invent — made innovations that could compute faster, can do calculations faster. We could add, but they could add a lotmore, faster than we could.

But the AI does something that we didn't originally think that a machine could do better — it could learn faster, learn within a well-defined area. And it — well-defined, we have to be able to train it, there's a whole set of circumstances where we know how to have the AI learn. It can learn how to play games like Go, but actually solve lots of very difficult things that would take us years and years — maybe decades — to learn.

So that ability to learn, combined with an ability to process and store information, has given us the ability to solve problems that were inconceivable a short while ago. You know, right — an example we're all experiencing right now is the quick identification of the pathogens that caused Covid-19 and the development of the mRNA vaccines. That's an example of the advances in biology. These complex molecules, we would never have been able to study, learn about them without AI and the computing technologies that we have.

So in many areas of science today AI is playing a very critical role. And in many ways AI — the speed of AI may be particularly relevant as we enter a world in which changes occur more rapidly, partly because of technological change. But there is an interesting dynamic here. AI causes faster change and we need AI then to cope with the faster change that AI is in fact causing.

MR. KORINEK: You emphasized before that there are several ways in which things can go wrong that move us away from the optimist scenario. And the first way was still under the assumption that markets are working well, but you observed that there can be redistributions. What do you view as the particular risks of AI in that context?

MR. STIGLITZ: Well, first, you know, we have to recognize it's very hard to imagine a world with perfect markets, because one of the things about perfect markets is we would all have insurance against the changes that AI might bring about. And then there

wouldn't be the redistributive consequences because we would have been insured against those redistributive consequences. But we don't have that kind of insurance and it's almost inconceivable for some of the reasons that I've explicated in my work in asymmetric and imperfect information. Those insurance markets just don't exist.

But obviously in the absence of those insurance markets, innovation, AI, can lead to a decrease in the well-being of workers, especially unskilled workers, workers whose jobs are more routine, because AI and robotization more generally has the effect of replacing workers. So one way of thinking about is for those — you know, those who are used to thinking of demand and supply, we can create artificially more humans, or more machines that act like humans. And in that way we increase the supply of human-like services and that would diminish the relative return of humans and at least lower their relative position in our society.

And so that's the dramatic — the concern about what is called — what Hicks called labor-saving innovation, innovations which at current market wages and prices, interest rates and so forth, reduce the demand for labor. And it reduces the demand for labor at current wages, that mean to clear the labor market wages will have to go down.

And equally, even if you don't — and I don't believe the labor markets are competitive, and I don't believe it — it even further weakens the bargaining power of workers because now workers have to compete against machines. Machines don't join unions, they don't give you trouble, they don't go on strike. They do sometimes break down, but they don't get Covid-19. And there are many advantages that machines have.

So it changes the bargaining power of workers, and especially workers for whom AI is an effective replacement. And that creates greater inequality and greater inequality has all kinds of societal consequences of the kind you hinted at at the beginning of our conversation.

MR. KORINEK: That's a really powerful way of describing the labor market effects of AI.

Now, if we add market imperfections to the mix, you have shown in some of your work that innovation may even reduce output and make everybody worse off in the economy. Can you explain how that could happen?

MR. STIGLITZ: Well, there's some obvious examples where in economies with market imperfections, innovations are directed at exploiting those imperfections. Obvious example, we have a lot of innovations that increase — that have the effect of increasing market power, increasing monopoly power. And we see that in some of the platforms.

Now, we know that in order for markets to generate efficient outcomes, we have to have competition. But if innovations undermine competition, then they're undermining the bases on which the market economy functions and the basis on which we have confidence that the market economy is efficient. Example, going back to AI, one of the things that AI has done has been to — the platforms get an enormous amount of data. The few platforms then have a competitive advantage over those who don't have that data. More data is the key input into AI. And that means it accelerates, it amplifies market power.

But then they use that market power not only maybe for efficiency, but also for exploitation. One form of exploitation is being able to have targeted pricing. So technically economists say they have the ability to extract consumer surplus. We used to think about those firms that are more profitable are delivering better goods. That may not be the case. The firms that are more profitable may be the firms that are better able to exploit market power, to amplify market power, to exploit individual vulnerabilities, to target individuals and extract the consumer surplus. So the old theory was that innovation gave rise to more profits because it made goods that were cheaper or more to the liking of individuals. Now we realize that innovation can give rise to more profits by enabling more exploitation of a whole variety of kinds.

Now, that's the most obvious example. But then there are some examples that again are fairly familiar to many individuals where you see the ability of AI and

innovations to explicitly exploit market imperfections in ways that aggravate them rather than remedy them. One example, for instance, we have regulations, regulations to protect us against when we go into a taxi, regulations, you know, hotels to make sure that they're safe and clean and, you know —- we have a whole set of regulations that make our society function better. But one of the things that some of the platforms have done has been to make profits out of regulatory arbitrage. In other words, getting around the regulations by circumventing —- you know, Air BNB may circumvent the regulations designed for hotels, Uber may circumvent regulations that are designed for taxis. And the profits may be partly due to a service that was not better provided, increase in efficiency. But part of the profits comes from regulatory arbitrage.

Another example that I've written about a while ago, is we have regulation against front running in financial markets. But there's been a lot of writing about flash trading as being an innovation. The new technologies where you have trading in nanoseconds that have no social benefit that anybody has been able to discover, but is a 21st century version of front running with all the disadvantages, all the undermining of the informational role of markets.

So these are all examples where actually markets are weakened, the efficiency of markets are weakened.

Two more examples and then I'll stop. You know, we're used to thinking one of the reasons that markets are efficient in general is that everybody faces the same prices. I don't know if you remember back to your basic undergraduate economics course, it was the fact that everybody faces the same prices, that ensures that the marginal cost and marginal benefits of everybody are the same. And that's the condition that generates the efficiency of the market. The first welfare cure. Well, I had mentioned before, the ability with AI to target —- charge different individuals different prices. What does that mean? It totally undermines the basis of the efficiency of the market economy. So the irony here is that this innovation actually destroys the theorem on which rely for the efficiency of the market

economy.

Now, there's one more example I want to give, somewhat more complicated, and I don't know if I'm going to be able to explain it. But it goes back to the —— what is called the theory of the second best. Take an economy which is competitive, that works pretty well, except, for instance, there are not a full set of risk markets. The theory of the second best says improvements in one market without solving the other problems may actually lead us away from —— may actually make welfare lower, may make everybody worse off. And the example that I analyzed almost 40 years ago was one where we had two islands that couldn't trade with each other because the cost of transportation was too great. And it was risk and the risks were negatively correlated, and there's no risk market. And then we have an innovation that allows them to trade. We reduced the cost of transportation. And now the result of that is they can trade with each other. That has the effect with a negative correlation and output of stabilizing prices. But because there are imperfect risk markets, stabilizing prices with variability and output increases the variability in income. And the net result of that was that everybody on both islands was worse off. Now that was a stylized example, but it brings home the point that when —— the world we live in, we're never going to get to a perfect market, we're never going to eliminate all the market imperfections. And creating some innovations which would in a perfect market have been welfare increasing, may in an imperfect market be welfare decreasing.

And, again, let me give one practical example. Structured finance, many people though big move in the right direction. In fact, many people said structured finance is a step in creating what was called the Arrow-Debreu markets, completing the markets, making risks markets better. Not perfect, but making them better and therefore welfare enhancing. We all know that in the presence of imperfect information, imperfect risk market, structured finance helped create the financial crisis of 2008, which such devastating global consequences. I don't think anybody really thinks it was a welfare enhancement, once we take into account all the adverse effects.

MR. KORINEK: Thank you, Joe.

I think these were really a number of very good examples for this theory of the second best for this theory of how making things a little better in one market may actually lead to welfare losses.

Now, when technological innovation makes some in society worse off, and you have already emphasized it, people are not well insured against these downsides of the technological progress, then we'll have greater need for social protection, for social insurance. But at the same time, the losers may also lose some of their political bargaining power. How concerned should we be that there could be this type of political economy amplification effect that as some people lose out from progress they also lose some of their bargaining power and the insurance institutions that are meant to protect them are in fact abolished?

MR. STIGLITZ: Yeah, I am very worried. And it was one of the central points in my book, "*The Price of Inequality*".

Basically, there's a nexus where you have economic inequality can and often does give rise to political inequality and then political inequality gives rise to more economic inequality. Whether it does or not really depends on the political institutions. But if you have a democracy like the United States with Citizens United where you have a set of revisions that allow money to have a disproportional role in politics, then if there is more income inequality, those with more income and wealth may use that wealth to campaign for initiatives for a legislative framework that will serve their interests.

And this can take two forms. They can argue that we don't want to have low taxes and not engage in redistribution, not have social protection, and they argue against social protection. But more broadly, they have an incentive for having a weak state, a limited collective action. Why? Because with strong collective action you are more likely not only to engage in the kind of redistribution or social protection that I just described, but you may engage in market regulation. Market regulation includes stopping exploitation of market

power, individual vulnerabilities, all the market imperfections. And when you look at the sources of wealth in a country like the United States, one recognizes that a very large fraction of those at the top have gotten a significant —- not necessarily all, but a significant fraction of their income and wealth from one form or another of what we economists call rent-seeking, for market exploitation, for exploiting vulnerabilities, from exploiting one form or another of market imperfection.

And so they want to maintain that. You know, example that we all see. The coal oil companies want to be able to continue to exploit the environment even though we all know it imposes enormous social costs on the rest of us and on our future generations. We know that there's a social cost of carbon, we know that the social cost of carbon conservatively is $100-$125-$150 a ton. But we can't get the legislation to implement that because the vested interests are adamant about maintaining their wealth that they derive from exploiting the environment and thereby exploiting all the rest of us.

Now, I want to say that this is not inevitable. You know, it is something —- you asked me am I worried, and I am worried. There are some countries which at least in the past have taken another logic to this. Some of the Scandinavian countries, small, said we are small countries, we have to be open to globalization and innovation. And they've said we are committed democracies. How do you square the circle of being a committed democracy and committed to innovation and globalization —- which they have to be given their size —- knowing the potential adverse distribution effects? Well, the answer is we have to make sure that the vast majority of our citizens benefit from innovation and globalization. And so that's why they set up a framework of shared prosperity. So they view that as a necessity given their commitment to democracy.

Now, what worries me now, and going back to your first question and first remarks in the very beginning, that if innovation is not benefitting a majority and one wants to be able to maintain the political power to continue with inequality and exploitation, there's only one answer. And that is leave democracy, undermine democracy. You can't square

the circle. And unfortunately, it seems to me there is a lot of worry that is the direction that some part of the United States is taking.

MR. KORINEK: That's indeed a very worrying concern.

Now, let's try to be a little bit more optimistic. And we have both made this case in one of our joint papers, that we should attempt to steer technological progress, that policy should take a more active role in affecting what direction technological progress should take. Can you tell us a little bit more about that? And in particular, how do these ideas also apply to artificial intelligence?

MR. STIGLITZ: So first let me try to explain a little bit more about the idea that there are different kinds of innovations. I talked before about innovations that created robots or AI that essentially replaced labor, what Hicks called labor-saving innovation, innovations that lowered the demand for labor at existing wages. In the beginning of the 20th century we had a lot of innovations that increased the demand for labor at existing wages and those were labor-using, not labor-saving innovations. They drive up the wages and that was part of the reason we have that kind of shared prosperity. Interestingly, in the last part of the 20th century, beginning of the 21st century, that historical experience led many economists to think that innovation automatically would benefit everybody. You know, a rising tide lifts all boats. But there was no inevitability of that. It was a particular historical episode. We're not going through another historical episode where that may not be —– there's strong reasons for that —– to believe that is not going up.

Conceptually, one way of thinking about the distinction is as little twist in language that I sometimes use. We talk about artificial intelligence, AI, and the other one I sometimes talk about is IA, intelligence assisting innovation. And an example of that is the telescope or the microscope. Our eyes could only see certain things and then we had these instruments that enabled us to see more. We became better humans, if you want to think about it that way. We became more powerful in what we could do.

So that's just an example of how innovations can actually make us more

productive. And that's the example where it could increase wages.

Now, there are many things that affect the direction of innovation. We could be using our scarce resources, research resources, to figure out how to replace unskilled labor, make machines that are like automatic tellers. Or we can use our scarce resources to discover better ways of making renewable energy. You know, these are both research projects. Economics is about resource allocation. And we could allocation our research teams to solving one problem or another. All of us as researchers decide what to work on, and we could work on one problem or another. I could have been a physicist and I debated whether — I almost was. And that was a decision. And so this is a matter of resource allocation.

And the first insight that I want to emphasize is there is no theorem that says markets solve the problem if allocating research resources efficiently, that the market solution to the pace and direction of innovation has no optimality properties. I've already hinted at that — what drives innovation is profits, and there is not a clear link between those kinds of profits and social well-being. So that suggests that we ought to be steering innovation.

In particular, in one of my earlier papers I show that there was a bias towards labor replacing innovation. So there was a bias on the private side to reduce labor costs and create as a result more unemployment. Now, what we seen in the last ten years as society, as individuals get more focused on climate change, that we've had enormous successes in, for instance, reducing the cost of renewable energy. It's not an accident. Most of those innovations have not depended on discoveries that have been made in the last ten years. Most of those discoveries, or at least a large fraction of the discoveries, could have been made 20 or 30 years ago. Why didn't we make it? Because we weren't thinking about that issue. We didn't have the market signals, it wasn't viewed as a societal priority. To me, one of the most optimistic — you said let's have some optimism here — is how far we've gone in reducing the cost of renewable energy in the face of very little pushing on the

part of government. Very weak policy. And yet we've achieved enormous — well, there has been a lot of societal attention. Young people see the issue of saving planet as an absolutely existential issue. And so it's nice that so many people devoted that kind of energy. But we've had enormous success with just a little bit of steering.

And that leaves me to say well what would happen if we had more steering. Now, what are the tools that we have? Well, there are a number of tools. One of the basic tools is basic research is basically supported by government. The mRNA vaccine was ready to go in a sense because we had invested as a society to basic research in developing the platform, the idea. And we could be spending a lot more of our basic research and ideas that would save the planet rather than save unskilled labor.

So public investment is the first pillar. The second one is pricing. Some of the pricing is quite obvious. We don't price carbon. I mentioned the social cost of carbon. And if we don't price carbon, we have no incentive to — we have limited incentive to innovate. And so having a price of carbon would divert more of the AI and other kinds of research resources towards saving the planet rather than creating more unemployment.

Regulations have some of the same role. If we say you're going to have — you have to meet these targets, like we did in the automobile. That's like Congress would say creating a shadow price, but that's a peculiar language that economists use. The fact is it directs research to meeting that goal. And an enormous amount of success in doing that.

But we also have to recognize that unintentionally, for the most part, we've been steering innovation in the wrong way. So examples are the Federal Reserve went in and sets the interest rate at zero. What is it doing? It's saying the cost of capital is zero. It's not the scarcity value of capital, but it's the cost of capital. So a firm deciding whether it should do innovation to save capital or to save labor is encouraged to save labor. So we are encouraging labor-replacing innovation by our monetary policy.

At the other side, the fact that we impose so many taxes on labor and we give preferential treatment to capital gains and to all kinds of capital taxation, does exactly

the same thing. So many of our tax policies are discriminating — you know, we have always talked about how they hurt labor relative to capital in a static context. In this context of steering innovation, they are also encouraging innovation to save labor, to create more unemployment, rather than to save capital, which would — the effect of which would be with less capital required per unit of output, there would be more capital investment that would increase productivity and that would increase wages.

MR. KORINEK: Those are great examples and I hope that we will find more ways to inspire and encourage our AI researchers and developers and entrepreneurs to engage in more innovations and directions that use labor and that bring us together.

Now, let me turn to another core expertise of yours, Joe. You made some of your greatest research contributions in the economics of information. In fact, that's what you earned the 2001 Nobel Prize for. And I should say I still remember how you dutifully showed up to teach our first-year graduate course that very morning when you got the call from the Nobel Committee.

But, yes, speaking about information economics, the traditional focus has been on incomplete information. But as we observed before, advances in digital technology and AI have really led to a rise in misinformation rather than just asymmetric information, and particularly in the context of social networks. How do you view this development as an expert in the economics of information? And perhaps the most difficult question, what can we do to better equip our society to deal with this?

MR. STIGLITZ: Well, I am very worried. Analytically, most of my research was based on the predicate that there were asymmetries of information, but we had a legal frameworks that meant that when you disclosed information it was fruitful. We've now moved into a world where there is extensive mis and disinformation. Now, in some of my earlier work, I actually explored the incentives for mis and disinformation. I wrote a paper with my colleague Bruce Greenwald on why it is that we need to have fraud laws to stop mis and disinformation. We all know that we want good laws, truth in advertising. You know, so

it's been in the back of our mind that the market on its own isn't able to solve these information problems and we need legal framework to encourage truth.

As another example, we know that disinformation can do harm and that's why we have libel laws that allow for a suit on torts associated with mis and disinformation.

One of the mistakes that we made — understandable at the time — was to exempt the platforms from intermediary liability. That was the whole section 230. And this has unleashed a level of mis- and disinformation that was unfathomable, partly because — and here is where AI plays a rather negative role — AI is able to identify who would be receptive to different messages. And so what we have now is a market, which is very non transparent because no one knows who is getting what message. And the business model is engagement — an underlying engagement is enragement and underlying enragement is polarization. So the business model of AI in the United States has been polarization. They made money out of polarizing our society.

Now, it hasn't been as bad as it's been in some countries, like Myanmar, where Facebook has been central in creating genocide. So we fortunately don't have that, but it played a role in the insurrection on January 6, they played a negative role in a whole range of digital harms. There are things that we can do about this and still be consistent with our commitments to free speech, free press, our first amendment. We've always recognized that we need to balance those protections with societal harms. We have restrictions on child pornography, crying fire in a crowded theater. So we know that we may need to rebalance those concerns in light of the enormous social harms. And there was never a guarantee on virality. That was never a part of our First Amendment rights. Now, of course, you might say it was a flaw of the formers of the Constitution not to think about virality, but that of course was not conceivable. And so that's an example of where original intent is absurd in reading into the Constitution.

Interestingly, Europe is now making great progress in thinking through these issue. Better than we've been doing. And I know that Macron under the French presidency,

the EU hopes that they will pass two pieces of legislation, the Digital Marketing Act and the Digital Services Act, both to ensure more competition, but also to protect against these digital harms.

MR. KORINEK: What you are describing are really quite massive externalities that AI is creating for our society and presumably, since there are externalities that are quite inefficient.

MR. STIGLITZ: Exactly.

MR. KORINEK: Now, we have been talking about the redistributions generated by AI. Now, one concern that I'm particularly worried about is the implications of these redistributions for developing countries. What do you view as the main concerns there and what do you think is it that advanced countries, like the U.S., could do to help?

MR. STIGLITZ: Well, I am very concerned. We talked about before how in principle advances in technology, including AI, could make everybody better off, produce more output with the same input, better products, so forth. And what we said though is the market may make some group worse off, but we could — not necessarily will — we could redistribute from those who are the gainers to those who are the losers to make everybody better off.

But when the gainers are in one country and the losers are in another country, that kind of redistribution is much more difficult, maybe impossible. And that is the concern about developing countries and AI because —– and innovation more generally. Innovation in its current form is having the effect of saving on unskilled labor. We had talked about labor replacing. What is the labor replacing routine workers, what is the asset that is abundant in developing countries? It's raw labor, it's unskilled labor. And so you're decreasing the value of the asset that is the core asset of developing countries. And many of the innovations are going to be also saving on resources. We're committed to addressing the problem of climate change. There are a large fraction of the developing countries that are deriving much of their income from oil, gas, and other fossil fuels.

Now, there are some countries, like Bolivia, that may benefit — as they lose from gas, they may benefit from lithium. So there are going to be some materials that they'll gain from, but there will be — so this is an example of the distributive effects can be very complex, but there will be some big losers. And that can have — that will have very big global implications.

Let me just think what will happen in those parts of the world, the Middle East, that depend on natural resources or those parts of the world in which unskilled labor is the main asset that they have. I just worry about the wages being pushed down to subsistence levels, unemployment, and political and social unrest, migration pressures. And it is hard to believe that these inequalities won't have consequences beyond the borders of these countries.

Now, what can the advanced countries do? Well, there are two things that they can do. One of them harks back to what we were talking about a minute ago, steering innovation. We can try to steer innovation in ways that do not harm them as much, do not undermine — you know, that may enable them actually to become more productive with their endowments of unskilled labor so that would actually be enhancing the returns to the very low unskilled labor that they have.

The other thing, obviously, we can do is some forms of assistance, of a variety of kind. One of the tools we have is the special drawing rights of the IMF, kind of IMF money that we can grant — essentially no cost to the American taxpayers or taxpayers in other countries. We did that in the midst of the pandemic, $650 billion. My view is that we ought to do it on an annual basis and target it toward the developing countries and emerging markets with a commitment that the rules of the game of the SDRs as they go in proportion to the quotas. That means a large fraction majority go to countries that don't need it, but they could recycle it to the countries that do need it in ways that would enhance their development.

MR. KORINEK: Thank you, Joe.

We have almost arrived at the end of our fireside chat, but let me ask you one final question. We have emphasized that it's not a given that progress will proceed in the right direction and how high the stakes involved are, and so on. So I wanted to ask you do you have one actionable piece of advice for our audience, or perhaps one for policy makers and one for innovators on what they can tangibly do to steer technological progress in a desirable direction.

MR. STIGLITZ: Well, I think that typically innovators don't think through the social consequences of their innovation. And the big advantage of a market economy and the basic welfare theorems is normally we say you don't have to do that. Just look at what maximizes your profit. Adam Smith's theorem was that if you do what is best for you, do what is best for society. In this area it is clearly not true.

So both policy makers and innovators, firms, need to think a little bit more deeply about the full consequences of what they're doing and ask the question, if everybody acted like I do, the golden rule, what would our society be like.

MR. KORINEK: Thank you, Joe. And thank you for joining us today and for inspiring generations of economists and policy makers and technologists with the work that you are doing.

And as you have been listening to our conversation today, I hope that you have also found inspiration and actionable insights so that all of us together can work towards steering technological progress in a direction that benefits everyone and that strengthens our society and reflects our values.

I hope that you'll join us again for future events in this series.

And now let me give a round of virtual applause to Joe.

Thank you again.

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

Carleton J. Anderson, III

(Signature and Seal on File)

Notary Public in and for the Commonwealth of Virginia

Commission No. 351998

Expires: November 30, 2024