Talent Versus Luck

Or, 99% Of Everything Is Rubbish And That's Why Stupid People Are Rich

Introduction

A paper by Pluchino, Biondo & Rapisadra 2018 (PBR18) claimed that luck dominates over talent in making people wealthy. Their central claim is that wealth follows a power law distribution whereas abilities are Gaussian, and they use numerical, agent-based models to demonstrate that luck can explain this. I was interested in but skeptical of this result. For a start, the observational evidence is very limited. While it seems well-established that the distribution of wealth follows a power law, with the richest fraction having far more wealth than their proportion of the population (of course the exact discrepancy is controversial), whether this reflects luck or some other factor is much harder to judge. PBR18 claim that abilities are known to follow Gaussian distributions, but give no evidence to support this (they cite only a single research proposal). Although the difference in wealth and talent distributions *may* indeed be interesting, there are some much more fundamental questions we should ask first.

How does one even measure ability, anyway ? Clearly people's ability to earn money follows a power law distribution, at any rate, but what of their more fundamental skills ? I believe I.Q. tests show Gaussian distributions but I don't know about results of other intelligence tests (e.g. the Berlin numeracy test; feel free to suggest others I could look up)[1]. While it might be possible to estimate genuine ability for some manual tasks (e.g. running speed, amount of coal shovelled per hour etc.), most performance depends on *quality* as well as *quantity*. This is much harder to estimate. Few would argue that the number of citations a scientific paper receives is a perfect measure of its true quality; artistic skill is clearly subjective; even technical works such as engineering and medicine have an associated quality factor that's hard to quantify. And even when ability can be measured, skill and motivation are different - as are their causes. Intelligence and personality seem to be at least heavily influenced by genetics as well as education, so rewards based on ability would still be (in part) due to the luck of the genetic draw. Michael Young pointed out that even if you could reward based on some Platonic ideal measurement of ability, this would be giving a significant unearned advantage to the children of the - ahem - lucky recipient. So determining a fair system of reward is far from trivial.

Measuring success is another can of worms I don't really want to open. PBR18 use the terms success and wealth interchangeably. Wealth may be difficult to measure, but it's definitely easier than success, which is a much more subjective and ambiguous concept… one man's paradise is another man's hell. Therefore I will only ever use the term "wealth" hereafter unless I mean to be deliberately vague.

Suppose for argument's sake that there really is a discrepancy between the distributions of talent and wealth. This is *consistent* with talent not being the dominant factor in accumulating wealth, but it is not *evidence* for this because it doesn't automatically render any other explanations less plausible. It could be that a small increase in talent reaps a disproportionately larger reward; salaries may simply follow a power law but still be based on ability - in which case the distribution of talent is essentially irrelevant. While not disputing the many examples PBR18 rightly cite in support of luck playing a bigger role than we might want to acknowledge, it may be a step too far to claim that luck is the dominant factor at work.

[1] Though if I recall correctly, I.Q. does correlate with "success", in some sense, but only to a point.

Nevertheless, the PBR18 numerical model is very interesting and worth exploring further. The model consists of a virtual 2D world populated with agents and events. Agents represent people, who each have some starting wealth, a fixed amount of talent, and are placed at random, fixed positions. Events move through the world in a random walk. Whenever an event and an agent closely intersect, the agent's wealth is affected. If the event is a positive one, then there is a chance the agent will double their wealth. This chance is proportionate to their talent, i.e. more talented agents have a higher chance to exploit lucky opportunities. In contrast, if the event is unlucky, then the agent halves their wealth regardless of their talent. A summary of the properties from the PBR18 standard run (the fiducial run hereafter) is shown in the table below.

Property	**Value**
No. agents	1000
No. events	500
World size	201 x 201
Timstep duration	0.5
No. timesteps	80
Lucky event fraction	50%
Talent distribution type	Gaussian (truncated range [0:1])
Mean talent	0.6
Standard deviation of talent	0.1
Wealth distribution	Uniform
Mean wealth	10.0
Standard distribution of wealth	0.0
Agent movement speed (units/timestep)	0.0
Event movement speed (units/timestep)	2.0
Event movement direction	Random
Radial proximity to trigger events	1.0
Talent affects luck status ?	No
Talent affects lucky events ?	Yes
Talent affects unlucky events ?	No

PBR18 find that this luck-driven model reproduces a power law wealth distribution with a slope of -1.33. Thus it successfully transforms a Gaussian distribution of talent into a power law of wealth. Additionally, they find that the 20% richest agents possess 80% of the total wealth (Pareto's rule, though I'm not sure that this is the case in real societies). However, there are two obvious problems :

1) The model gives luck an enormous influence. Chance dictates whether an agent encounters an event, chance dictates the luck status of that event; chance plays a significant (though not the only) role in determining whether the agent can exploit it. There's very little way in which talent can manifest itself in this model.
2) Wealth can only double or halve. There's no debt, no spending, no salaries, no wealth transfer between individuals. This would seem to make the development of a power law inevitable.

Thus the model would seem designed to succeed. It seems interesting to ask how far this model can be pushed before it breaks. Of course, there are many (!) other simplifications of the model that could be addressed, but I would like to try and deal with the model on its own terms as much as possible.

Reproducing the PBR18 model

The first task in exploring the model is to reproduce it. Any efforts to improve or modify it must first demonstrate they can successfully reproduce the same basic results as the original, otherwise the comparison is unfair. To do this, I wrote a Python code based on the description in PBR18. The code is open source and can be run [online in a web browser](). It automatically produces the same plots as in PBR18 (below, plots from my code are on the left and those from PBR18 are on the right). Some of these (such as the talent distribution) are trivial so not worth reproducing here.

One of the main claims of PBR18 is that there's no correlation between wealth and talent :

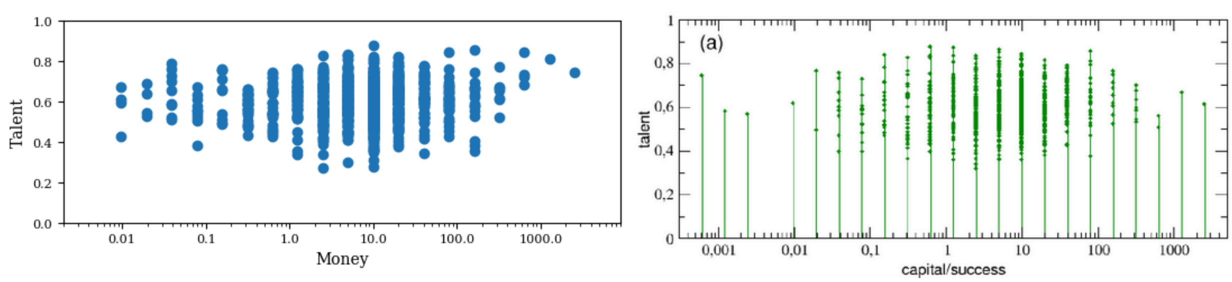

Note the logarithmic money axes. There is no correlation using either the new code (left) or the original (right). Later on we shall see why actually this is much more subtle than it appears.

Conversely, PBR18 claim that there is a very clear relation between wealth and luck, i.e. the number of lucky and/or unlucky events an agent experiences. First let's see the trend with respect to luck events[2] :

[2] I don't like some of these plots. Events are clearly the dependent variable so ought to be on the x-axis, but this would make it very difficult to compare to the PBR18 originals.

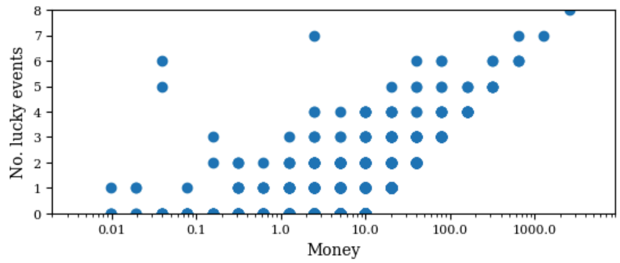

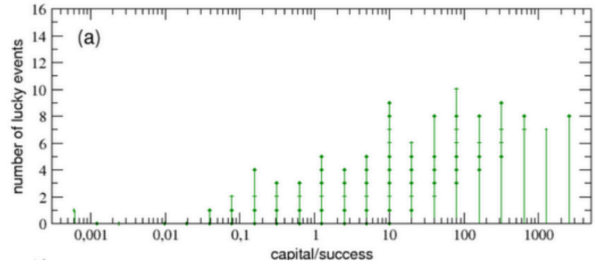

This claim clearly holds true (due to chance, there is a slight difference in the number of events that occurred in the different runs). The scatter is due to two factors : 1) agents can also experience unlucky events that decrease their wealth; 2) talent influences whether agents can actually benefit from lucky events. Again, we will consider the philosophical interpretations of this later on.

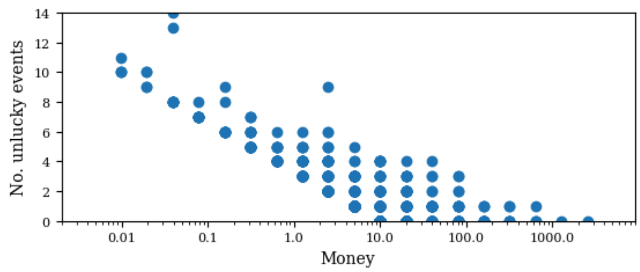

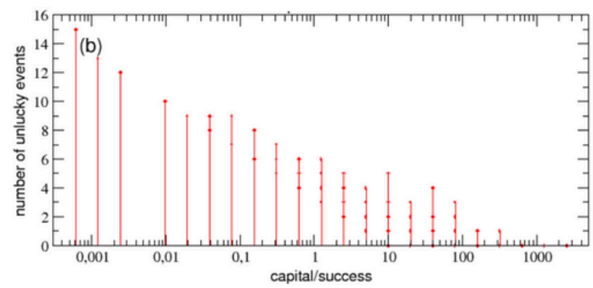

There is a very similar anti-correlation between the number of unlucky events and wealth, again with a slight scatter due to agents also being able to experience lucky events.

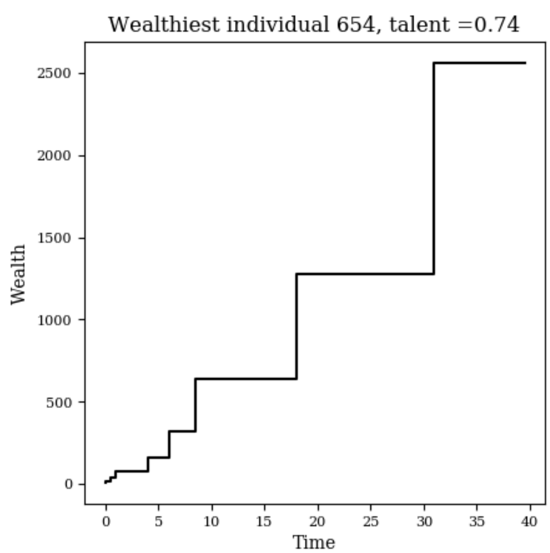

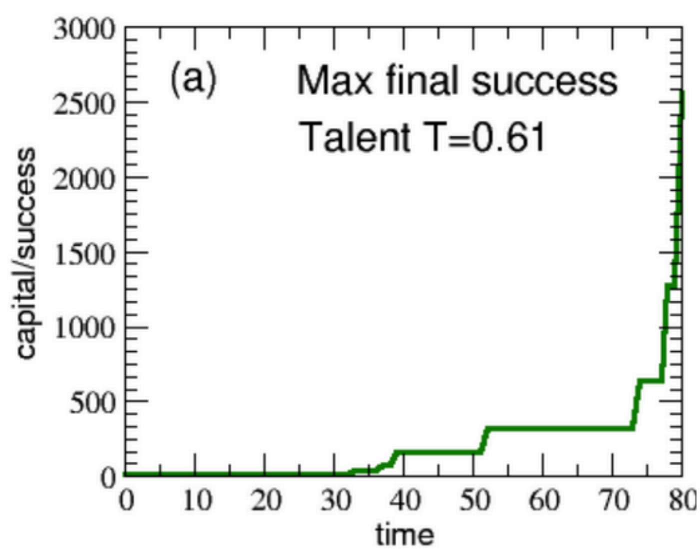

The evolution of wealth of the ultimately richest individual follows a broadly similar trajectory. In this case, in the new code the wealthiest individual did have talent a bit above average but in other runs this was not the case - it is a result of chance alone.

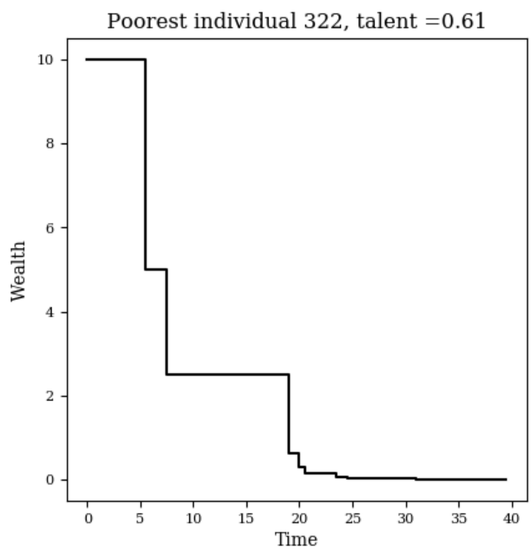

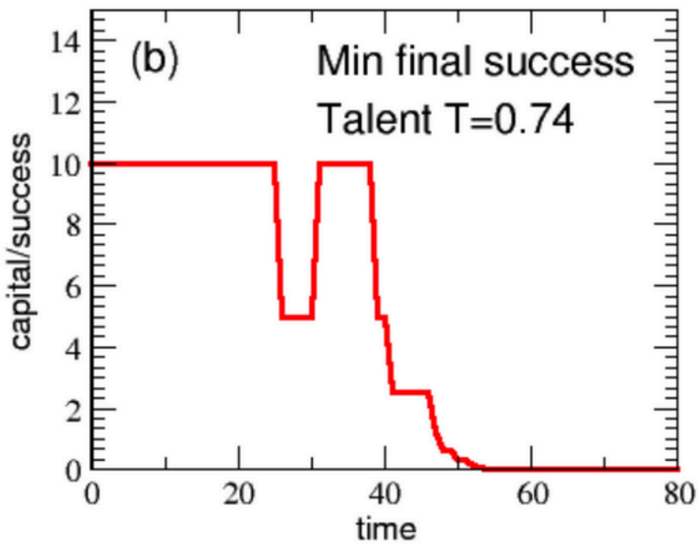

It’s pretty much the same story for the poorest individual. Note that the talent of the poorest individual in this particular case of the new code was slightly above average : many less talented individuals were more successful. So it appears that there’s no correlation between talent and wealth even at the extreme ends.

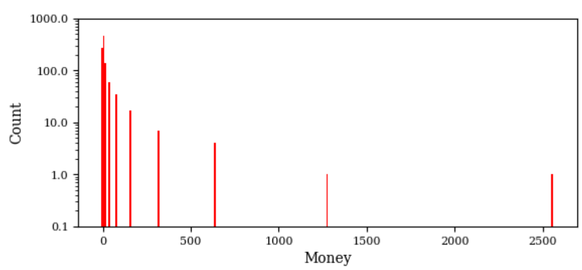

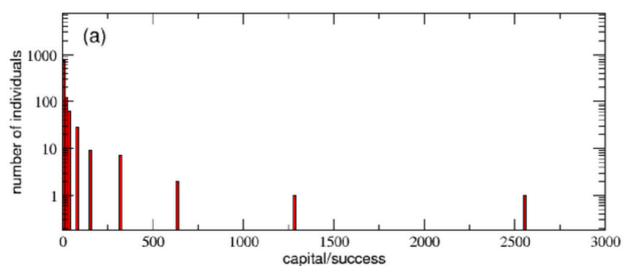

The distribution of wealth for the entire population is also very similar using the new and old codes. The slope of the power law is basically identical :

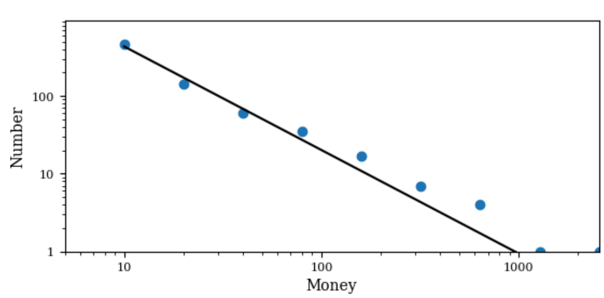

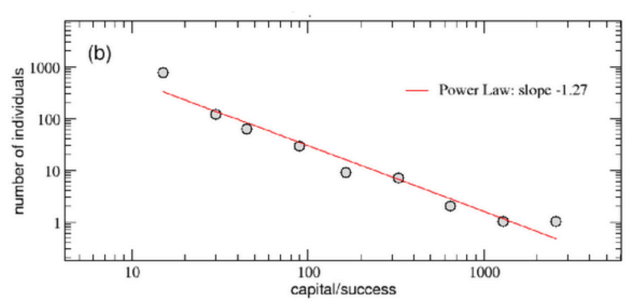

The black line in the left plot has a slope of -1.33 which PBR18 found from multiple runs; the slope of -1.27 on the right comes from a single PBR18 run.

Taken at face value, this all supports the basic PBR18 conclusion : that talent is far less important than wealth, and that luck is a valid mechanism to explain the difference between the (supposed) Gaussian distribution of abilities and (observed) power law distribution of wealth.

Another way of examining the wealth distribution is, as PBR state, to compare the wealth fraction owned by the 20% richest people over time.

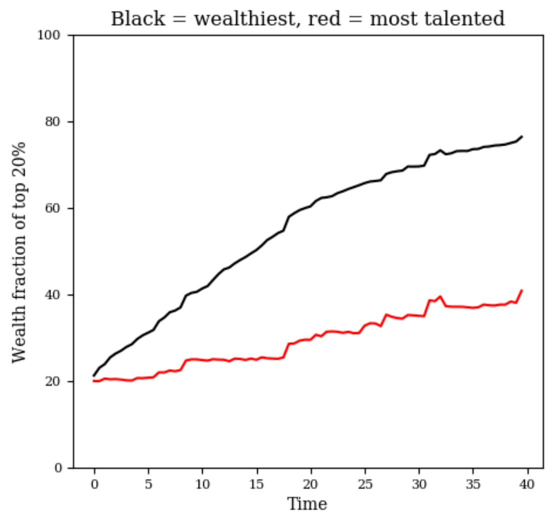

PBR18 do not show the time evolution, but state that the wealthiest 20% of the population own 80% of the wealth at the end of the simulation. The same is true for the new code. We can also see that the wealth fraction of the 20% most talented evolves relatively little over time (in some cases the curve is even flatter than this). Again, talent doesn't seem to count for very much. This "Pareto plot" is the single most useful diagnostic plot for analysing the simulations, although some trends can only be revealed with the others.

The agents in the Pareto plot are not necessarily the same from timestep to timestep : it is possible that an agent could have low wealth for most of the simulation, then suddenly experience a rapid succession of events and quickly become the richest of all. This raises the question of whether, in this model, the rich always get richer and the poor always get poorer, or if in fact wealth is more transient. We can answer this by finding the richest agents at the end of the simulation and plotting their entire history of total wealth. We see that here the curve is almost always rising : the rich *almost* always get richer (the same is true if we plot wealth fraction in the same way). Interestingly the wealth of the 20% most talented agents also shows a slow but steady increase.

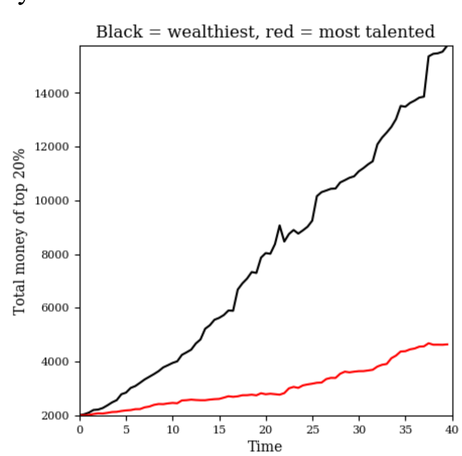

How robust are the conclusions ?

The model has a lot of different parameters and is unconstrained by observations. This makes it natural to ask how finely-tuned the model is : if, say, we doubled the number agents or allowed both the agents and events to move, would anything fundamental change ? It seems like a good idea to keep the model itself basically the same and only alter certain parameters before we try anything more drastic.

Unfortunately there is no single perfect diagnostic plot we can examine. The Pareto plot is probably the simplest though, so let's try that. Here's what happens if we try and alter a bunch of parameters. In each case, all parameters are as in the fiducial run except for that which is specified. It turns out the model is pretty robust to some things : changing some parameters makes very little difference. In all plots, the solid lines indicate the 20% most wealthy/talented while the faint lines represent the least wealthy/talented.

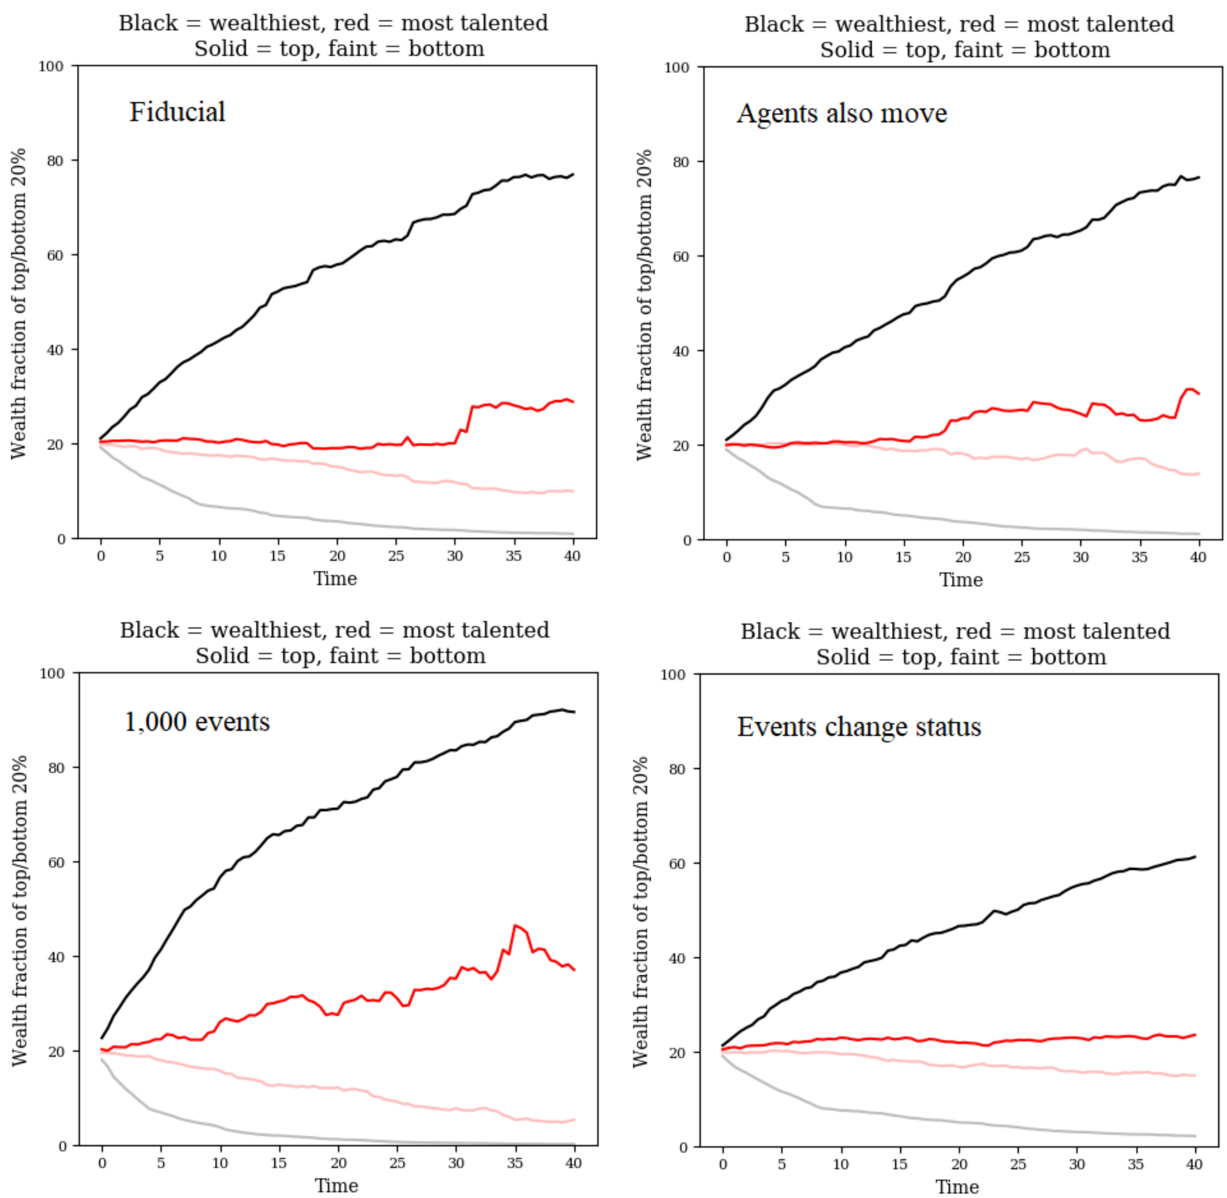

It seems to make no difference at all if we allow the agents to move as well as the events (both using a random walk of the same step size). There might be a bit more variation in the most talented agents for some reason (though this variation is quite normal anyway), but that's about it.

Doubling the number of events does alter things a little bit : the rich get even richer, and the most talented maybe get a little bit richer. This makes sense - more events, more opportunities to exploit. Not really surprising.

For no particular reason, I also tried making all of the events swap their lucky/unlucky status every time they move. This slightly reduces the wealth of the top 20% but that's about it. So you might think that event position makes no real difference, but that would be a huge mistake, as we'll see later.

There are other parameters we can alter which also seem by themselves to make no difference, but these are better explored as a narrative. This is not how I came to realise what was going on, but it tells a nice story so let's go with it. It turns out there's much more to the fiducial setup than meets the eye.

If we have 2,000 agents (twice the fiducial) we see that the most talented people now get a larger share of the wealth :

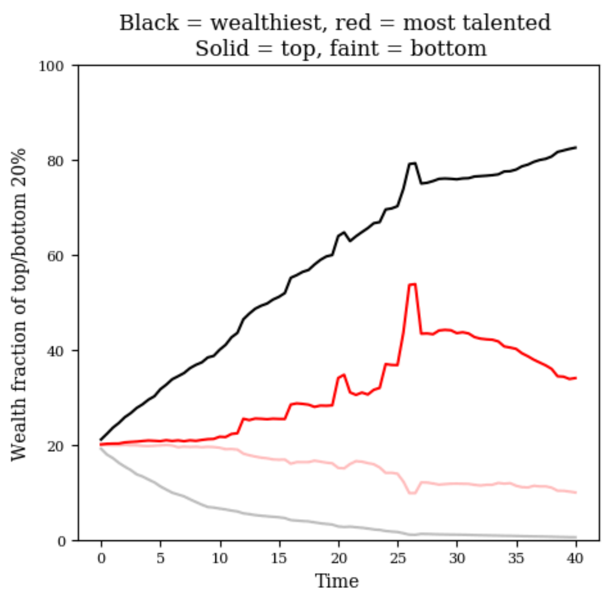

Why's that ? One reason is the Gaussian distribution of talent. With only 1,000 agents there are very few really talented people; with 2,000 there are significantly more. And they are much better at exploiting lucky events.

You make your own luck

But raw numbers aside, there are two interesting things to vary : talent and the events that it experiences. Let's begin with talent. As mentioned, the Gaussian distribution is a difficulty. What if we simply made the distribution uniform ? Keeping everything else constant (with the standard 1,000 agents) we get :

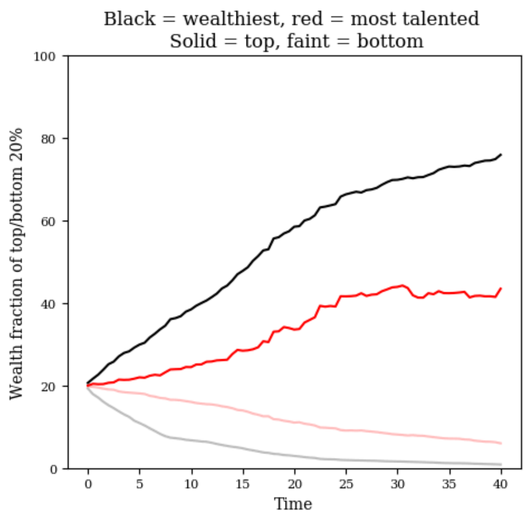

That’s markedly different to what we’ve seen previously. Those talented people and the richest people are no longer mutually exclusive categories. This is, after all, what we should expect : talent *does* play a role in exploiting opportunities, it’s just that thus far there haven’t been any serious levels of talent to explore. The Gaussian distribution of talent *may* be realistic (although that’s an open question), but if we want to investigate talent’s role, then this distribution is not appropriate because it greatly reduces the numbers of the most talented people. To have a fair comparison, we need to compare equal numbers of talented and stupid[3] people.

This explains the subtitle of the document. Since there are so many more stupid people than geniuses, it’s more likely that one of the mediocre will get lucky than one of the smarter people. But in equal numbers the geniuses *as a population* - there will always be unlucky exceptions - should *always* outperform idiots.

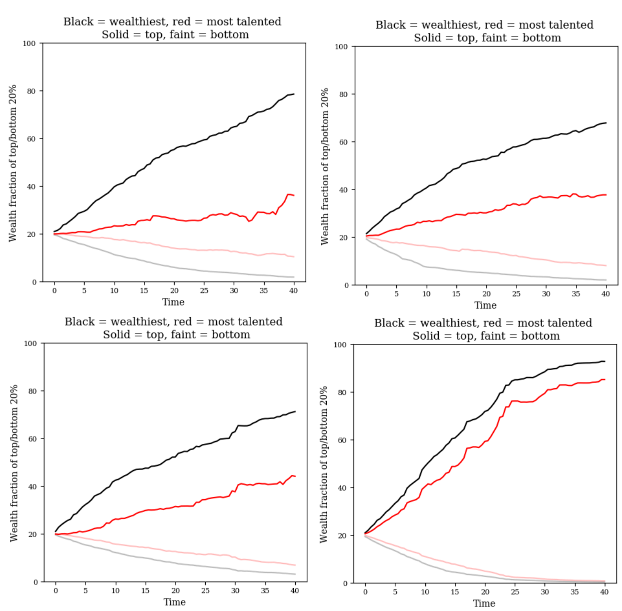

We can also alter talent in more fundamental and interesting ways, as on the left. Just as talent normally affects the chance for an agent to exploit a lucky event, so it can be set to affect their chance to avoid unlucky events in the same way (top left). This has very little effect. But if the luck status of an event is affected by the agent’s talent, then the more talented individuals begin to gain a greater share of the wealth (top right). Combining both effects gives a stronger result (lower left), and using a uniform distribution of talent as well leads to the most talented agents owning the majority of the total wealth in the simulation.

[3] It’s my simulation so I get to call them stupid if I want to.

Are these changes sensible ? It seems reasonable that there are different types of chance events. Some we have no control over, while the effect of others depends entirely on our own abilities. Winning a lottery ticket or inheriting money makes us richer regardless of any of our talents. In contrast, discovering a mineral deposit could make a geologist fabulously wealthy but have no effect at all on a lifestyle coach who has no clue as to what they've found. Similarly, some unfortunate events cannot be mitigated by talents (e.g. being robbed) but others can (e.g. most people who experience extreme suffering just suffer, but a few go on to win the Nobel peace prize). This is not a moral judgement on anyone; the best method to allow people to develop and exploit their talents is far beyond the scope of this document. The point is only that the effects of talent in the above figures seem reasonable. There does not seem to be any compelling reason to prefer any one of these models, including the PBR18 interpretation where talent played a much more limited role. It is likely that they all apply in the real world in different situations, but determining the relative frequencies of those situations is highly non trivial.

These charts also raise the question of whether we could now see a direct correlation between talent and wealth. Of course, we can answer that too.

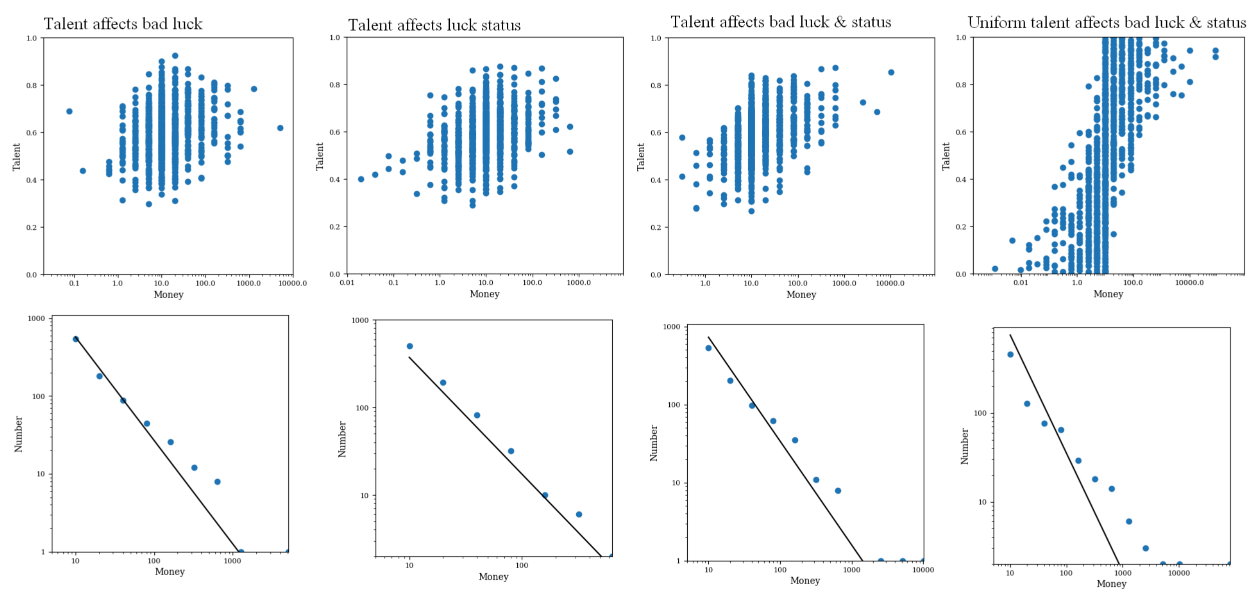

All of the alterations produce more of a wealth-talent correlation than the fiducial case, though the effect is much stronger if talent is allowed to affect the luck status of an event. We can also see that the uniform distribution of talent makes an even greater difference. The y-axis (talent) in the top row is the same for all cases; the Gaussian distribution clearly severely limits the upper and lower parts of the plot, disguising any trends.

In the bottom row we see that the first three cases follow the same power law as in the fiducial run (the black line has a slope of -1.33 in all cases - it's the variation in the axes that makes the slope appear to vary). But when we use a uniform distribution of talent we see a deviation. It still seems to be a power law, but with a flatter slope. It's the third plot from the left that's the most important : with a Gaussian distribution of talent, it's possible to produce both a clear talent/wealth correlation and a power law distribution of wealth. *Hence a difference in the distributions of abilities and wealth cannot be taken as evidence of the importance of luck.* And this model remains true, I think, to the spirit of the original.

In fact, even the PBR18 model shows evidence of this, though it's hidden in the noise. Since talent does influence the success of an event, there ought to be a talent-wealth correlation - maybe just a very weak one, but it should be there. The thing about the talent-wealth plots is that they compare **all** individuals. But as before, if we want to investigate the role of talent, we have to give it a fair chance. We saw how it wasn't fair to compare the success of a single genius with that of a hundred idiots, since sheer numbers give the idiots a much greater chance that more of them will be successful. Similarly, it's not really fair to compare individuals who happen to experience a hundred events with others who encounter one or two. By necessity, those who have more opportunities for success with become richer than those with greater abilities but no means to exploit them.

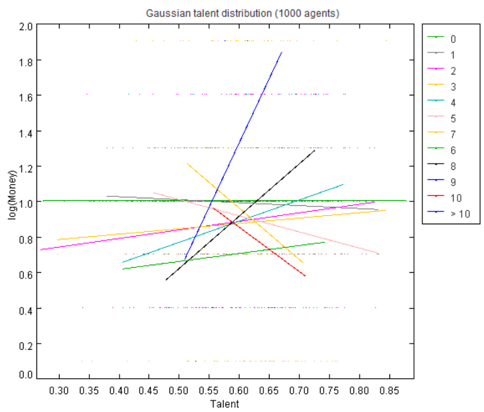

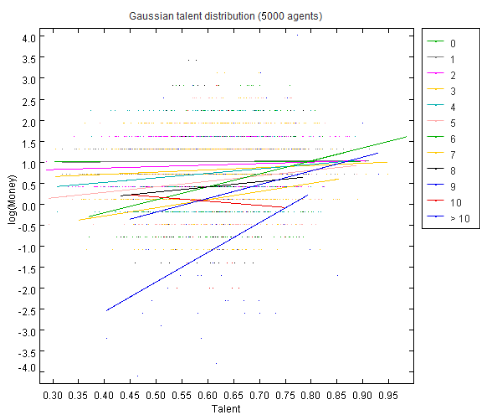

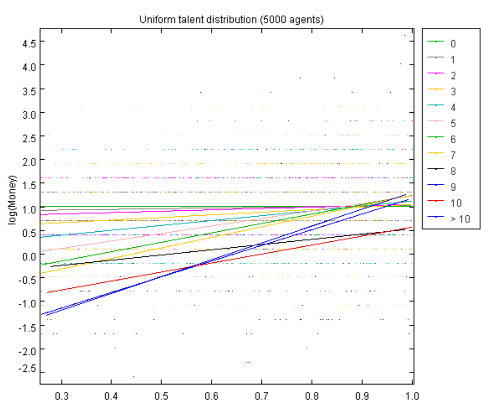

I apologise for the poor quality of these plots, I'll remake them with greater legibility at some point. What they show is the final wealth of every agent as a function of their talent (just like before only with the axes flipped, since talent is the independent variable), colour-coded by the number of events (of any status) they experience. The lines show linear regression fits to each agent group of equal event numbers. The first plot is the fiducial run. There is clearly little or no trend : the slope of the lines varies randomly. Increasing the number of events has no effect on the "trend"; it appears to be purely random. In other words, there is no trend in wealth as a function of talent regardless of how many events occur.

But once again the Gaussian distribution is a horrid thing. If we simply increase the number of agents, the best fit lines become clearly non-random. The more events agents experience, the steeper the slope : the greater the talent-luck correlation. It's a subtle result and hard to see without the best fit lines, but quite unmistakable. If we use a uniform distribution of talent (third plot), the trend becomes even clearer.

This doesn't mean that luck isn't the dominant factor at work : it clearly is in this case. What it demonstrates is the importance of considering the statistical effects. There's a trend here, subtle and hidden, but definitely present. Just as correlation doesn't equal causation, so too, in this case, does a *lack* of correlation not equal a lack of causation. If we had an equivalent observational data set, we'd have to be wary of interpreting a lack of talent/money trend as evidence that talent wasn't being rewarded. We would have to control for the number of opportunities for advancement being presented to different individuals.

A postcode lottery

The earlier tests seemed to indicate that the events were distributed pretty much randomly; it doesn't matter if the events change their status or if the agents are also in motion. This gave me the idea to make a parametric version of the code that could run several orders of magnitude faster than the standard approach. Normally, the distance between each agent and each event must be checked and every one evaluated to see if it's within range. But by first measuring the typical event occurrence frequency at each timestep (as well as its variation), it should be possible to select at random the appropriate number of agents and give those agents an event. It's possible to measure the rate of multiple encounters as well. The typical number of agents experiencing one or more events per timestep turns out to be ~40, so this is massively faster than the standard approach. The results were OK, but not quite the same as the original.

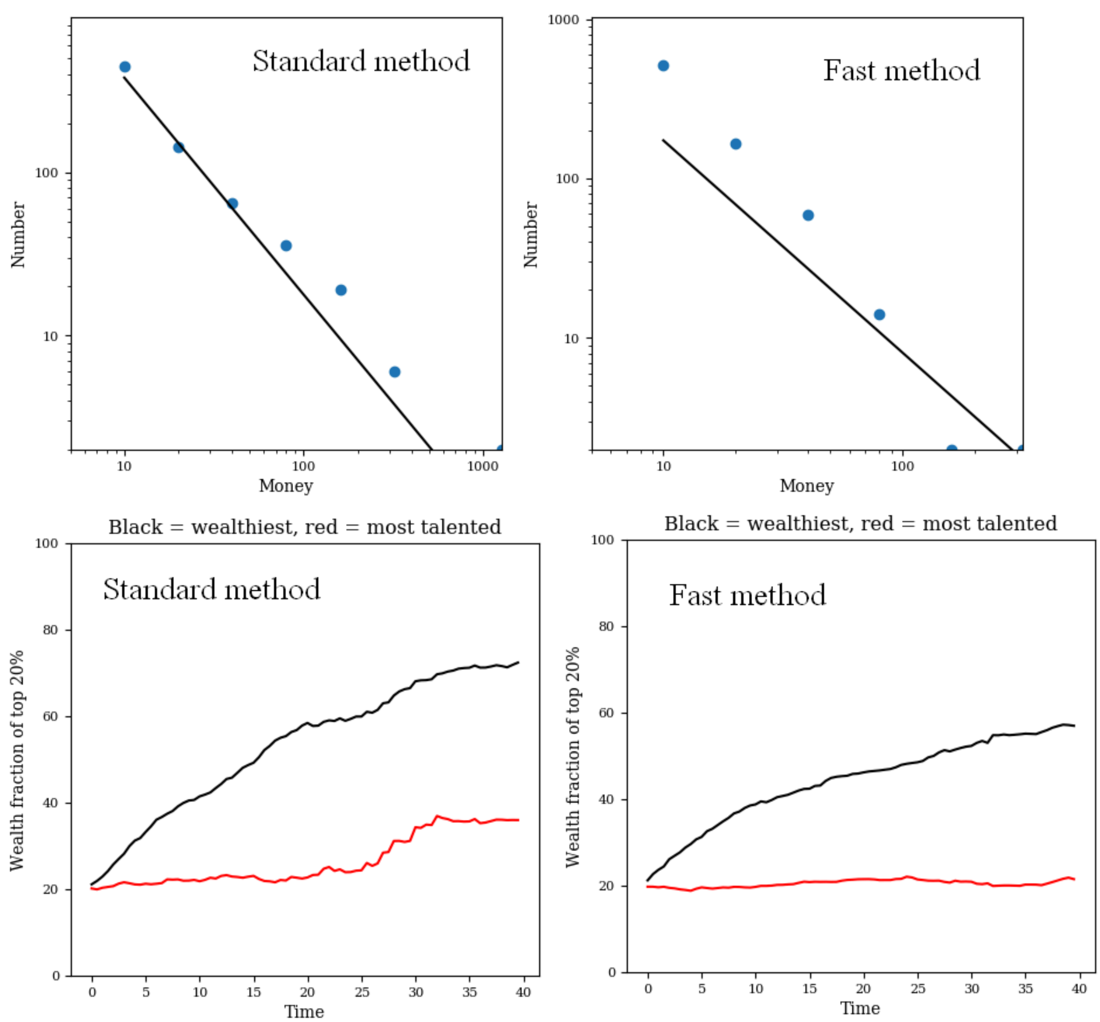

This doesn't look so bad, but even making very drastic changes (multiplying the event rate by factors of several didn't change the main result. Unlike the fiducial case, to alter the wealth distribution you have to make ***massive*** changes - or if you push it far enough, things just get plain ugly. Serious, Piers Morgan levels of ugly.

I’ll spare you the long tale of how I went about solving this. Suffice to say it was absolutely awful and I failed utterly. But I did learn something very interesting along the way.

At any given timestep, the events are distributed randomly. They are always distributed randomly in space… *but not in time*. With each timestep, they don’t move very far : their current position depends on their previous position. Yes, they move in random directions, but that doesn’t matter because they don’t move very far. What this means is that you sample the world at random, there are some areas where you’re far more likely to find an event than in others. This can be visualised by plotting the positions of events at all timesteps at once :

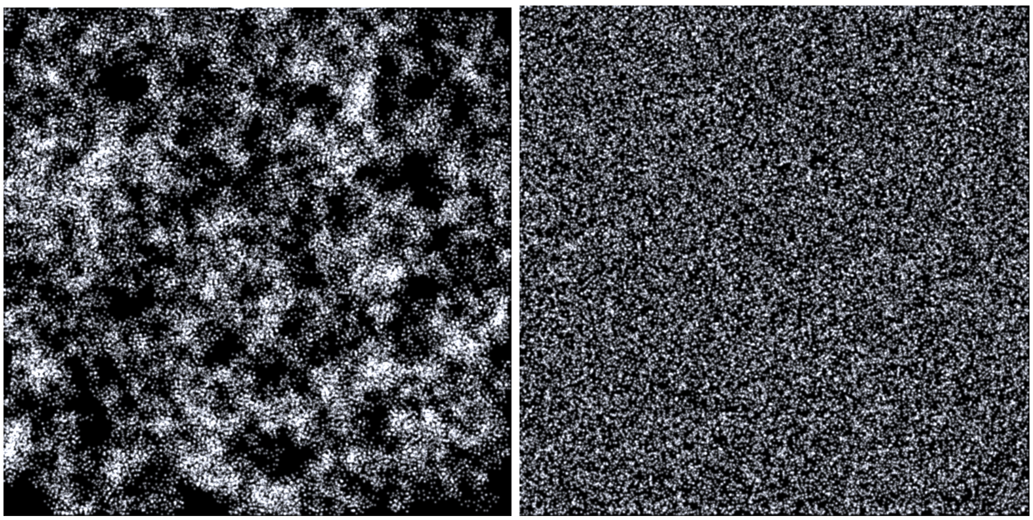

The plot on the left shows the random walk whereas the plot on the right shows the case of completely randomising the event positions at each timestep.

What I found was that my parametric method gave an identical Pareto plot to the case of randomising the positions each time. Actually, I’d already noticed that one simple way to change the wealth distribution was to change how far the events move. Below, the left plot shows what happens if the events move 4 units, which you’ll note is very similar indeed to the fast method shown in the previous plot. The right plot shows the case of the events moving only 1 unit.

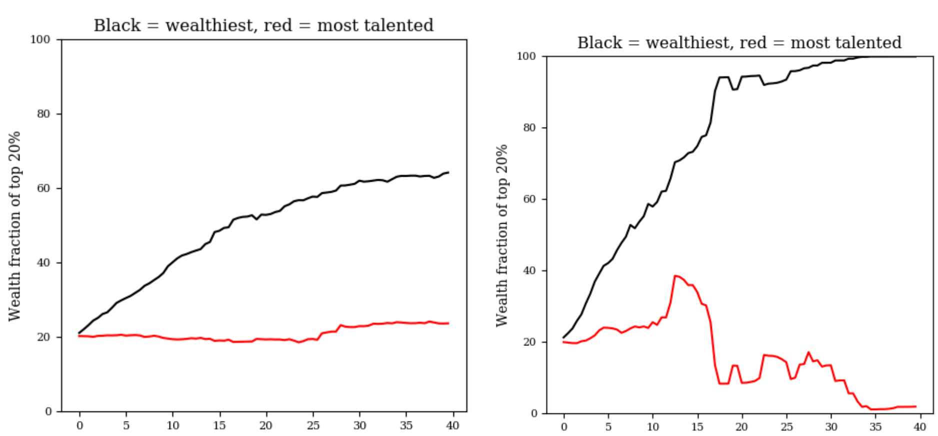

Choosing the events to move 2 units seems rather fine-tuned; it avoids the chaos (and total wealth domination of the elite) of them moving too slowly but isn't quite the case of truly, temporally random positions either. The fact that the final result is such a neat 80:20 rule in PBR18 is not an intrinsic feature of the model itself.

The above plot on the right is more interesting. What's happening is that agents are repeatedly encountering the same event multiple times in succession (an effect the parametric method cannot reproduce !). This can drive them to rapid wealth or poverty if they happen to start near a good or bad event. This can be manifested even more dramatically if we double the event effect radius (left) :

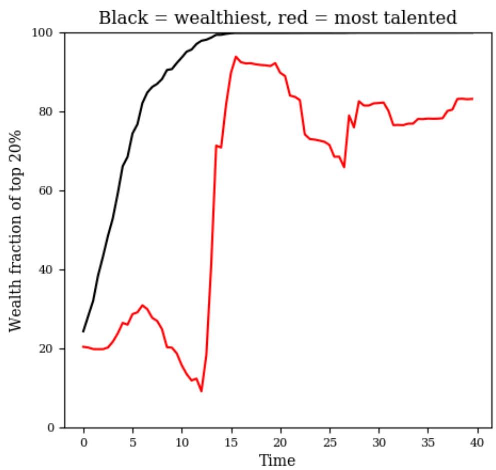

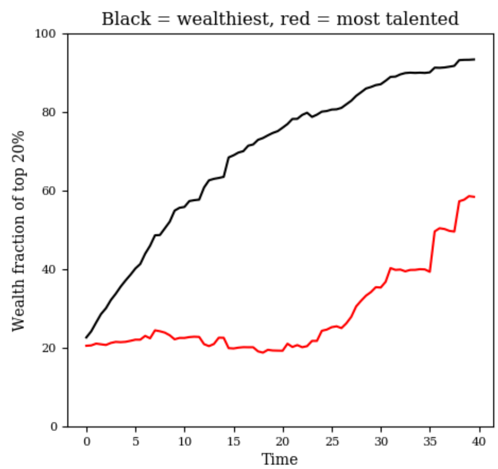

It gets pretty chaotic. There really is something like a postcode lottery at work here : opportunities and risks alike are not present uniformly but in distinct (though mobile) areas. This is very different from the case of simply doubling the number of events (right plot above). The chance of encountering a good *or* bad event doesn't really change much in that situation, in that the two will basically balance out, but if the effect radius is increased it makes the probability of multiple encounters with the same event much more likely. Although this does show that luck if a factor, it also shows the PBR18 model is very finely tuned to give the results it does. More interestingly, it demonstrates that the distribution of event positions is very important.

"I'm a celebrity, get me out of here !"

Thus far the simulations have used an initially uniform distribution of wealth, as PBR18 did, as this makes the effects of talent and luck easier to disentangle. But what about using a different distribution, say a power law ? That would seem to be a reasonable way to deal with the issue of inheritance (but if anyone's got any better ideas feel free to chime in).

The easiest way to code this is to let the simulation run once, then reset the timer to zero and move all the agents to different starting positions. This isn't the most efficient but it's slightly simpler than recording/creating a power law distribution and assigning agents to it at the start. What it results in is agents starting the simulation with a power law of money but nowhere near the events that led them to their initial fortunes (hence this is different from simply letting the simulation evolve for longer). Dunno how realistic that is, but it'll do.

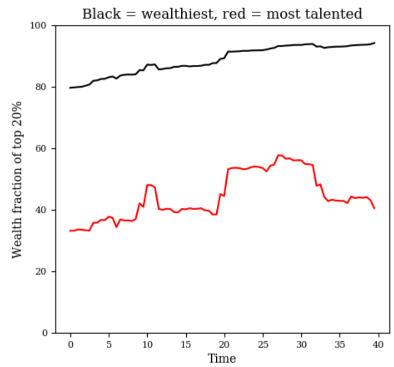

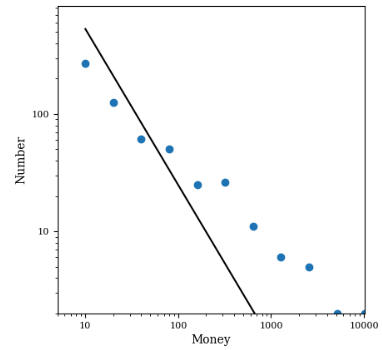

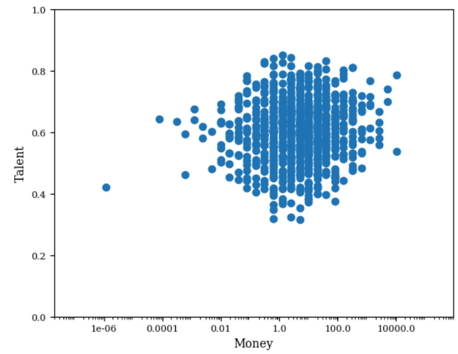

What we find is that the rich get richer and the power law flattens. There's still no obvious large-scale trend in talent/wealth. Interestingly, the fortunes of the most talented appear to be much more variable now. This is backed up in repeat runs, where sometimes the wealth of the most skilled can reach over 60% of the total but other times crash to a mere 3%.

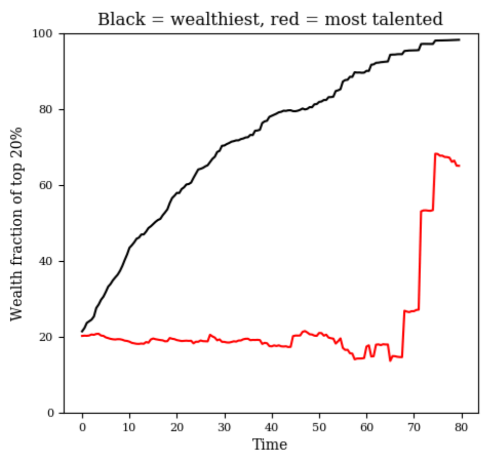

However, I don't think this is much a result of resetting the agent positions. If we let the fiducial run evolve for longer, we see similar behaviour (left). So it might be that, since the agent's movement speed means the encounters with events are not truly random, that rapid, successive encounters with good/bad events can happen : we just need to wait a little longer to see them.

Repeating the above method but allowing talent to affect both the unlucky events and the luck status of events, we get :

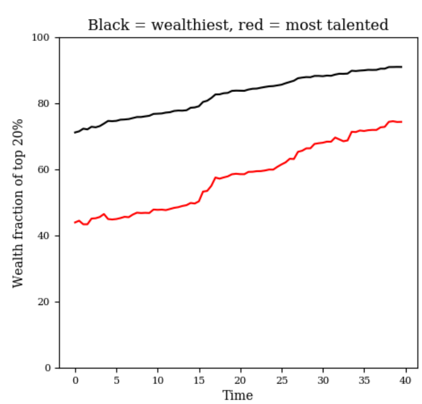

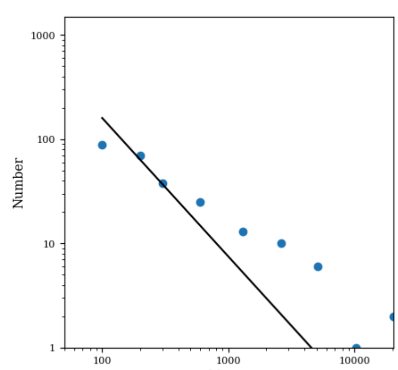

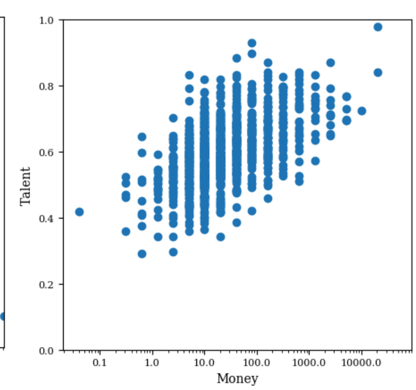

Though the Pareto plot can vary more dramatically. I suspect sometimes it's strongly dominated by highly deviant individuals. Something needs to be done about them. And it is, [later on](#).

For the sake of completeness, here's what happens if we do the same as above but with a uniform distribution of talent :

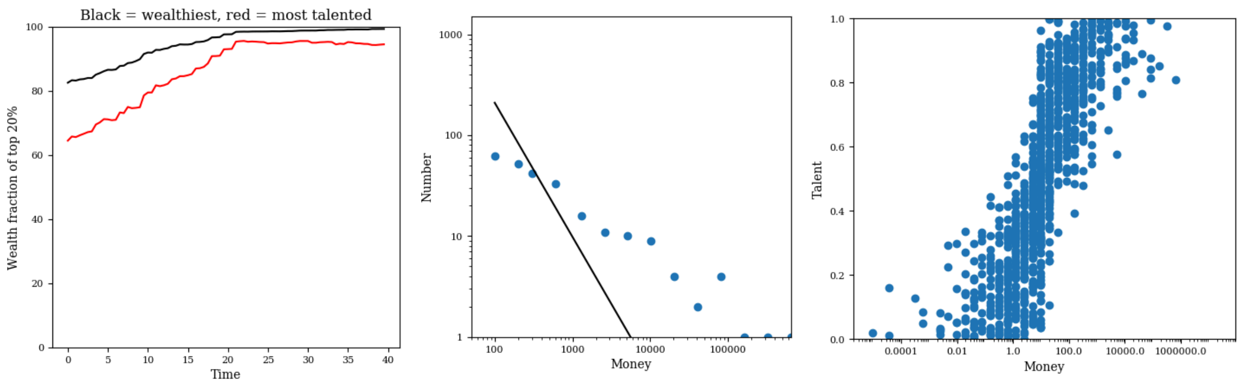

A bit more wobbly power law and the most talented become even more wealthy. In all cases, the rich inevitably get richer. Does that mean that initial wealth matters more than talent ? Probably. If you start with extreme wealth, it takes little to maintain it and a lot of serious errors to lose it, whereas if you start poor but talented, you'll have to climb an awfully long way to reach the top. Perhaps the issue here is that talent has no effect on the amount of wealth you can gain or lose. Well, we'll see about *that*.

One concern with the above is that the standard Pareto plots - unlike the absolute wealth evolution plots - find the most wealthy agents independently at every timestep. This makes it (somewhat) inevitable that the most wealthy will continue to get wealthy, almost (though not quite) by definition. A better analysis might be to see what happens to the agents who are the most wealthy at the mid-point of the simulation, at the moment their position is reset. The plots below show the same conditions as above but using the mid-point agents to find the wealthiest and poorest :

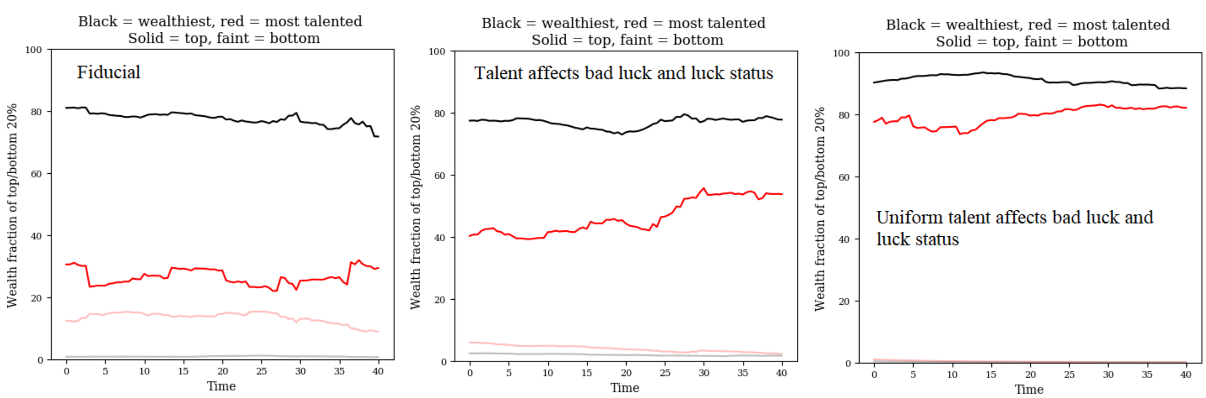

There is little change besides a flattening of the wealth of the 20% richest. Rather than continuing to amass an ever-greater fraction of their wealth, once placed in a new (but similar) environment they simply tread water. Their immense wealth prevents them from falling significantly, but they are unable to make any further progress.

A malevolent meritocracy ?

If you want an extended analysis of why a pure meritocracy is a terrible idea, consult Plato's *Republic*. We can't test anything remotely that sophisticated here, of course. But what we can do is make it so that talent has an even stronger effect. Instead of the simple doubling or halving of wealth, we can make the rewards and losses proportional to talent. This is not exactly a fair system. It could be interpreted as rewarding people differently for doing the same tasks, based on a perception of their supposed deeper underlying ability *without regard for their motivation*. That hardly seems like a sensible system : the point of a genuine meritocracy should surely be that rewards are only available to those of the appropriate ability level at all - they must still possess some drive to *actually* exploit those opportunities, not just be given money because they *potentially could* exploit them if they wanted to. If a professor of statistics decides to quit and open a newsagent, we don't expect to pay them more just because they have an unrelated qualification. We expect that they get a higher salary while doing the more difficult and demanding job, not to go on endlessly rewarding them for their existing awards. Similarly if someone with no qualifications makes an astute observation about some obscure aspect of statistics, we don't expect them to be given lesser credit than a professional researcher. And while people can't go bankrupt in these simulations, there's no lower limit on their poverty either, nor (yet) any way for their talents to vary[4]. It's not a very nice simulated world to live in, really.

On the other hand, another interpretation would be that more skilled people are able to more fully exploit opportunities that present themselves. A researcher who makes the same discovery as a layman is much more likely to understand its implications more rapidly : a slew of publications and grants may therefore follow, which is only fair if they do all the extra hard work that results. How you interpret these kinds of very simple simulations is highly ambiguous, which is either annoying or really interesting depending on how you look at it[5].

Anyway, I implemented this scenario in the following way. When an event occurs and affects an agent, the maximum reward is set in proportion to their talent. If their talent is zero, then their reward is to multiply their wealth by one, i.e. do nothing. If their talent is one, then their wealth is multiplied by two, and for everyone else there's a linear interpolation between these extremes :

Wealth = Wealth(Talent + 1.0)*

Similarly for unlucky events :

Wealth = Wealth/(Talent + 1.0)

[4] Nor is there any modelling of how the richest use their wealth, which is far beyond the scope of this investigation. It opens a very interesting can of worms though : how much talent is associated with benevolence, whether money/power corrupt etc.

[5] I think it's probably both. Applying the model to the real world would be completely ludicrous, but it's still interesting in its own right.

This means that geniuses (with talent = 1) always double their wealth when good events happen, whereas morons (with talent = 0) never benefit at all. Likewise geniuses always avoid bad effects completely, whereas idiots always suffer. Fair ? Hell no. But it’s an interesting idea.

Here’s what happens using the fiducial conditions but with the effects on wealth modified as described above :

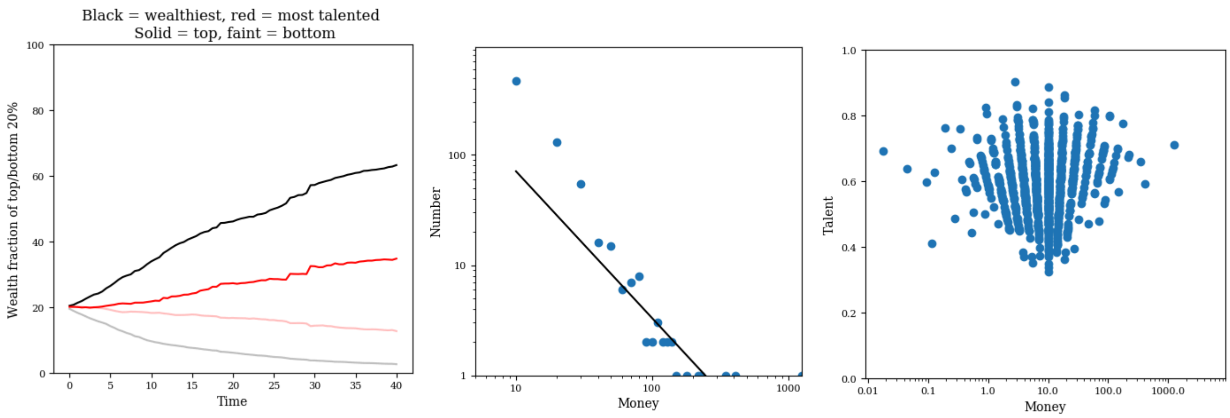

Not much change in the Pareto plot, except to lower the wealth of the richest slightly (and maybe increase the wealth of the most talented slightly). The wealth distribution changes substantially from the fiducial conditions - it’s still a (much steeper !) power law at low wealth levels, but with a long tail of a few much richer individuals. And the talent/money plot turns into a bizarre fan-shaped structure, which I wasn’t expecting at all.

We can repeat the now-standard procedures. Here’s what happens if we also allow talent to affect the success of bad events :

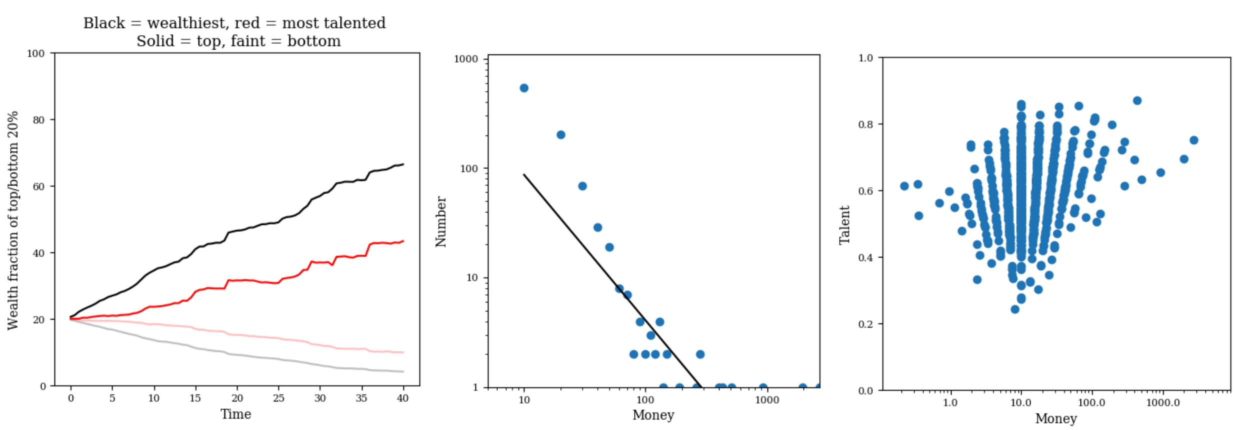

A somewhat different Pareto distribution from the fiducial run, and the overall distribution of wealth no longer looks quite like a power law (or at least has a lot more scatter). If we allow talent to effect the luck status as well as the success of good and bad events :

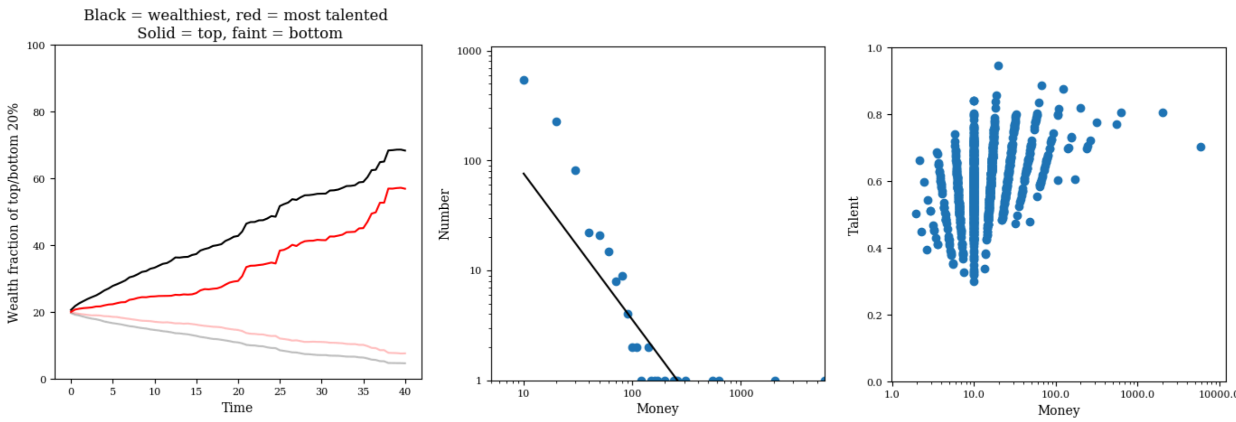

Clearly talent matters a great deal here. The power law again looks weird, though we should bear in mind the small number statistics at the low wealth end. Of course we can also try the same conditions but with a uniform distribution of talent :

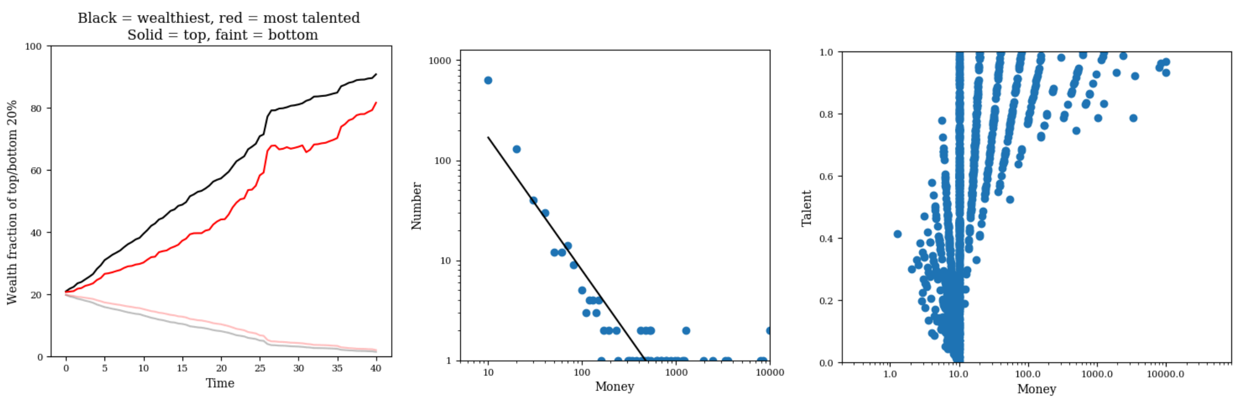

Which shows that the most talented are practically synonymous with the richest. The power law becomes clear again, albeit with a high scatter and a long tail where there are few members. And the talent-wealth plot is simply hilarious. However, the Pareto plot isn't massively different to what we saw under the same conditions but with talent not affecting the exact reward/loss level.

Safety nets and ceilings

Changing the initial wealth distribution seems to favour the talented somewhat, at least under the contrived conditions discussed above. Could another reason the most talented people don't make that much money (in the fiducial conditions) be due to sheer bad luck ? As PBR18 eloquently note : *"...if your life is as unlucky and poor of opportunities as that of the other agent, even a great talent becomes useless against the fury of misfortune."* What, therefore, would happen if we add in a social security system ? If we prevent agent's wealth from falling below a threshold, perhaps they can better endure runs of bad luck and so go on to eventually prosper. We could also consider capping the maximum wealth level, so that greater equality will make any trends easier to see.

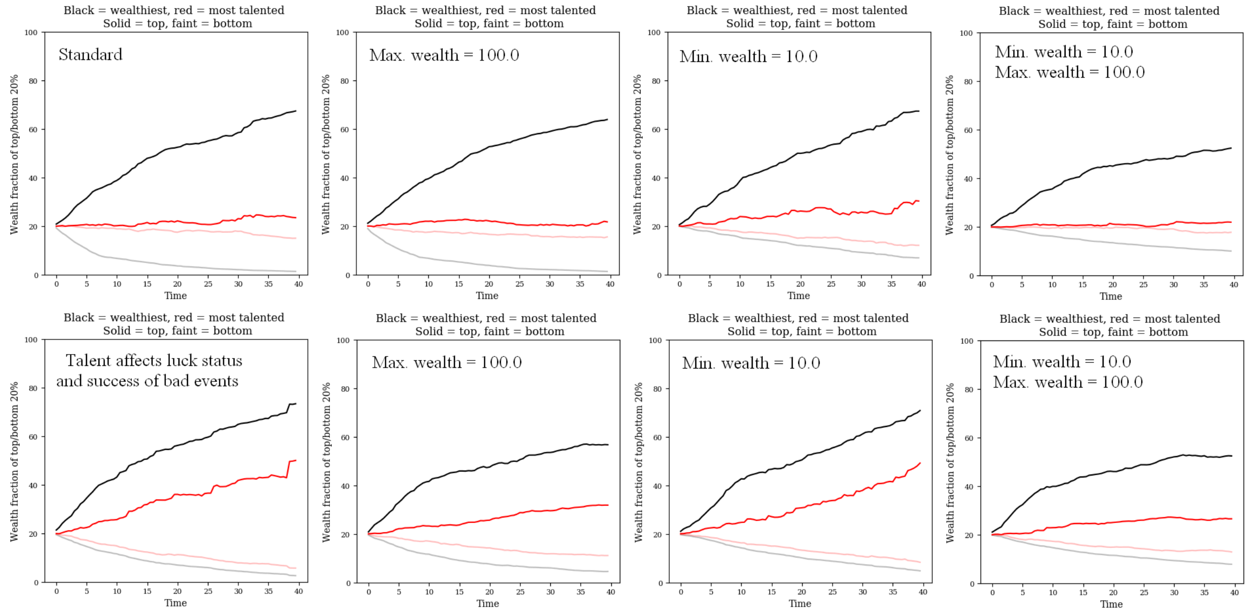

I also realised that I've hitherto been concentrating on the richest and most talented while neglecting the poorest and stupidest. So the above plots show the poorest people as a faint black line and the least talented as a faint red line [earlier plots were edited retroactively and didn't have these lines originally]. If we want a meritocracy, then presumably we want the most talented to also be the richest and the least talented to be the poorest. But what level of inequality we want is a much more difficult moral question. First let's look at what happens in the simulations.

The top row shows the fiducial conditions under the modifications of a lower and upper limit on wealth. The lower limit is referred to as a *safety net* and the upper limit as a *wealth cap*. Both are implemented in a very simple way : the safety net prevents wealth from falling below some threshold (here 10.0 units are used, the initial wealth level of all agents) while the wealth cap prevents them exceeding it.

The effects of the wealth cap can be a bit subtle. We've seen previously that while generally smooth and predictable, sometimes the evolution of wealth can be somewhat chaotic. The reason for this is that occasionally a few agents can have runs of lucky events in quick succession, driving their wealth far above that of the rest. This does not affect the talent-wealth diagram (which is logarithmic) but can significantly affect the Pareto plots. Thus a wealth cap can serve two purposes : if the value is reasonably high, it can prevent the existence of extreme outliers that skew the statistics; if it is lower, it explores the effect of wealth regulations. A value of 10,000.0 is sufficient to prevent extreme outliers for at least 400 timesteps without changing any of the above conclusions so we adopt this by default. To begin to see the effects of a wealth cap at a more fundamental level, from examination of the data a value of about 100 is required.

We can see that in the fiducial run, the rich are very rich (20% have 80% of the wealth), the poor are very poor (20% have < 1% of the wealth), and the most and least talented have about the same 20% fraction of the total wealth. Employing a maximum wealth cap makes very little difference to anything. However, a safety net has stronger effects. It doesn't do much at all to the wealth of the richest, nor help the fortunes of the most talented. What it does do is reduce the poverty level of the poorest and decrease the wealth of the stupidest agents, bringing those two demographics closer together - there is a 21% match between the lists of the poorest and stupidest agents.

Combining the safety net and wealth caps reduces the wealth share of the richest, but flattens and equalises the wealth of the most and least talented : talent matters even less than in the fiducial run. However there is greater protection for the poor, who now have 9% of the total wealth, with all curves essentially plateauing if we increase the duration of the simulation. The most meritocratic situation is to use a safety net; the most egalitarian approach is to use both a safety net and wealth cap.

In the second row of the above plots, talent is allowed to have a much stronger effect, i.e. affecting the luck status of events and granting the ability to avoid unlucky events. Without any restrictions on wealth we see (as before) a highly meritocratic result, with wealth of the most talented evolving in parallel to the richest (37% match between the two) and likewise, the wealth of the least talented falling in parallel to the poorest (41% agreement[6]).

The overall difference between the richest and poorest is rather extreme in this setup. Using a wealth cap flattens all the curves, though the agreement between the richest/most talented and poorest/stupidest is roughly the same as before : the society remains a meritocracy, but more financially equal. In contrast, a safety net doesn't do much at all. Hence the combination of the two effects is similar to, and dominated by, the effects of the wealth cap.

None of this should be taken as evidence for any particular social policy because that would be positively bonkers. Even the assumption that any of these particular scenarios is broadly applicable would be far too strong to allow us to state which approach would be better : there are simply too main unknown and unaccounted variables. The model doesn't account for the motivations and goals of the agents or how this is causally connected to money : it may be ideologically appealing to suggest that we need safety nets of caps, but neither can really be supported by this crude model.

We also have to return to the question of whether we really want a total meritocracy or not. Rewarding people for their work is related to but not the same as rewarding their ability : they must also possess drive and ambition, but we also don't want to reward cruelty[7] or give the least able so little that they are effectively crushed. Crucially, this model ignores the variation in talent and the highly complex feedback with wealth. Just giving someone more resources doesn't automatically give them more ability (except in the purely financial sense; it doesn't make them more skillful), but it does give them more opportunities to develop underlying abilities (purchasing books and educational courses, having sufficient time to be able learn new skills). Wealth can enable the development of greater abilities but it is also a result of greater abilities, while, at the same time, having the desire to use or develop those abilities is sometimes not connected to wealth at all. Thus, giving some people immense amounts of wealth might cause them to become extremely skilled at something whereas for others it might make them become, if anything, less skillful because it removes the need to *do* anything. It might, like other psychological factors, follow a non-linear behaviour - a little bit of pressure is an incentive to improve, but too much can be crushing; a

[6] Note these figures are somewhat uncertain and the code that measures this needs to be checked !

[7] Perhaps it's not talent we want to reward, but the more complex product of talent, ambition, and (for want of a better word) virtue. We want people to be rewarded for the work they've actually done, but not if they're liable to misuse those rewards and abuse others. Goodwill is always nice, but useless without the skill to use it; skill in a task is great but destructive if used malevolently. The combination is more than the sum of its parts.

bit of economic security can be beneficial and encouraging, but too much luxury may be corrupting. And, I'd guess, the precise way in which those feedback loops operate could vary strongly from person to person. The goal, presumably, is not just to reward people based on what they've done, but also to encourage them to do more. Unlimited rewards and unrestricted poverty don't seem like good ways to do that.

We can't possibly account for all this in the model. Perhaps if we had some much more limited, highly specific real-world scenario we could adapt it to genuinely reflect a real situation. But the general case is absurdly vague. Perhaps, though, we could try allowing talent to vary in a simple way…

There's no substitute for experience : a false meritocracy

Real people are at least capable of learning, even if they don't always enjoy it very much. Someone who's being doing a job a long time is almost always better than a complete novice. How does experience increase someone's talent ? I dunno. I would guess that varies both with the individual and the situation. Mistakes are sometimes said to be more informative than successes. Maybe that's true, maybe it isn't. The model lets us play with altering talent however we like. This makes the parameter space to explore potentially enormous. While this means deciding how to apply the model to the real world is even more difficult, it also reveals that even this simple two-parameter model is fraught with complexity. So perhaps we can be forgiven for not trying to draw conclusions about reality or make the simulation genuinely realistic just yet.

Perhaps the most obvious way to start is to use the fiducial conditions but allow talent to vary when events happen. In the plots below, talent is increased by 1σ whenever a good or bad event occurs. Previous inspection shows that the typical event frequency means that there shouldn't be too many agents for which this would result in a crazy increase in talent, though the talent range has been kept restricted to the standard [0..1.0]. Naively, one might expect that in the model it shouldn't matter whether we allow good or bad events to affect talent, since the model is incredibly simple. The number of good and bad events is the same, so while we expect the changing distribution of talent to have some influence on the outcome (knowing that there is an influence of talent in the fiducial conditions), whether the events which affect talent are good or bad should make no difference to the final result. Well, this turns *out* not to be the case.

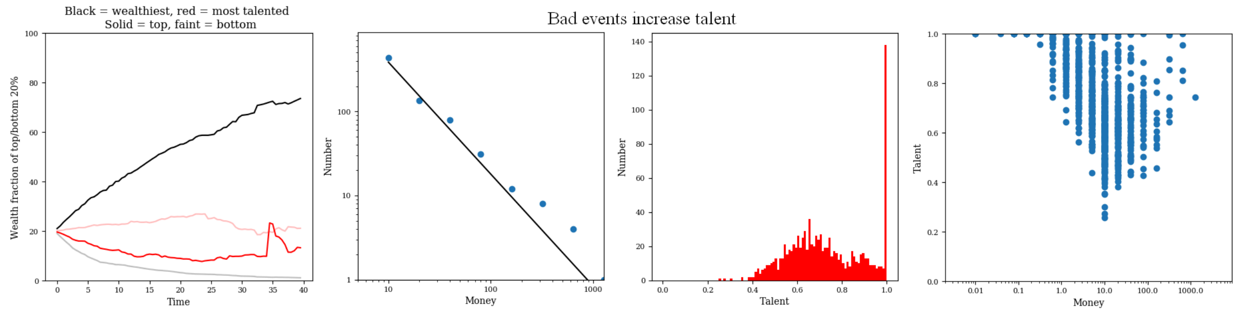

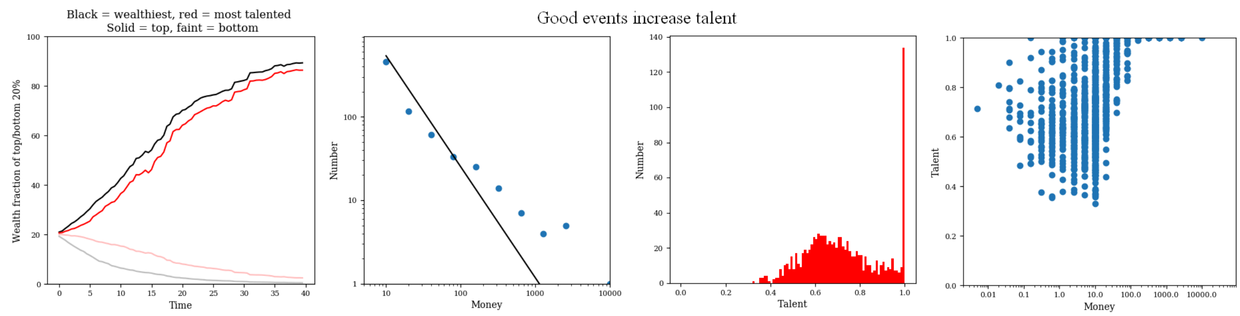

We see that the results are radically different depending on whether we allow the agents to "learn" from good or bad events. If bad events provide experience, the result is anti-meritocratic, with the most talented agents being among the poorest (though the stupidest aren't especially rich). There is a clear anti-correlation between talent and wealth. In complete contrast, if good events provide experience, then society becomes strongly meritocratic with an extreme difference between the wealthiest and the poorest and a clear, positive talent-wealth correlation.

This does *not* tell us something profound about society or the reward system we use. Rather it provides a nice example of correlation not equalling causation. Consider the case of improving talent in the event of a bad event. This means that as the agents get poorer, they increase their talent. But the effect of talent is very weak in this model (albeit partially only due to the small statistics), so more talented agents don't have a significantly greater chance of increasing their wealth. Essentially, the poor people keep getting smarter *but don't stop getting poorer*. The opposite is true for the agents who get smarter as they get richer.

With the previous models, we were trying to see if we could get wealth to be causally connected to talent. We wanted conditions under which society could become more meritocratic in our simple model. Here we have in effect done the opposite, using wealth to drive talent. In the case of rewarding success, it may appear that we have made society more meritocratic but there is no more a causal link between talent and success than there was in the fiducial run. We have, essentially, educated the already wealthy rather than educating people in order to increase their wealth. It's as though we decided to find a bunch of rich layabouts and force them to complete a full university or vocational course and then declare we'd got ourselves a meritocracy. It's all rather perverse.

We can see that if we use the model where talent affects bad events and the luck status, then we again get meritocracies regardless of which events we use to adjust talent :

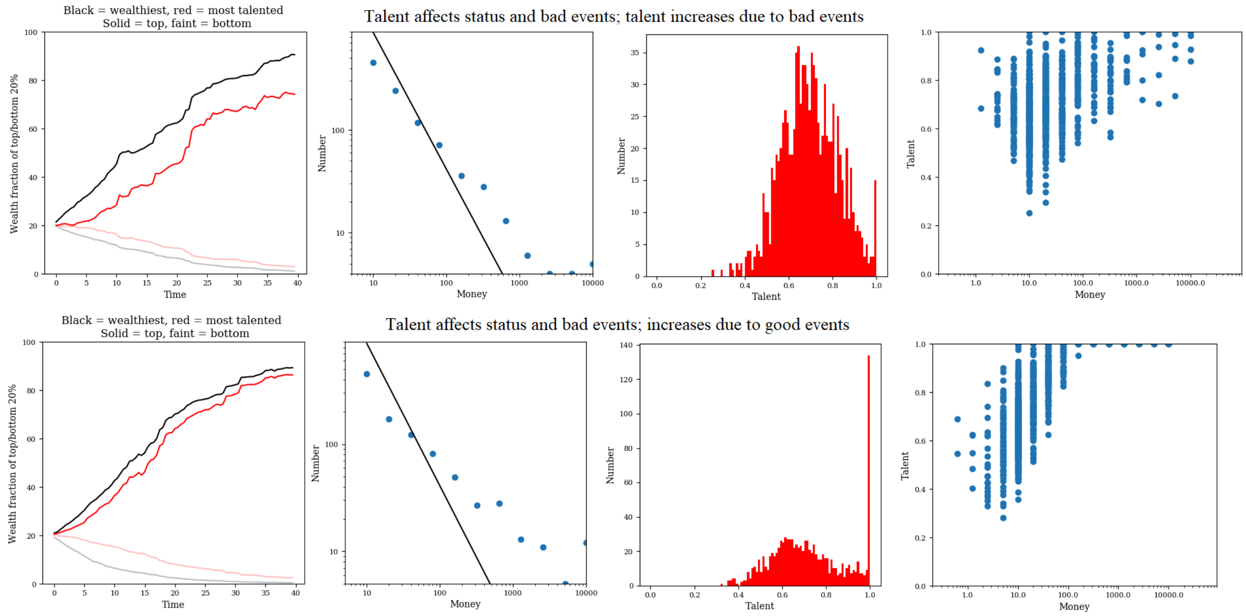

And we can use this to design systems which are even more extreme (see figure below) by allowing talent to decrease as well as increase. If we allow bad events (i.e. decreasing wealth) to reduce talent while good events (which increase wealth) increase talent, we get an uber-meritocratic society (or so it seems), whereas if we do the opposite we get even less correlation than in the fiducial case. Note that the upper figure uses the standard effects of talent, i.e. it's the ability to exploit fortunate opportunities but nothing else. Even with this highly restricted effect, we can produce a society in which the trends correspond to a perfect meritocracy so long as talent varies. Of course, we know this isn't a true meritocracy at all. Rather, it's more analogous to finding all the rich people and giving them some intensive training whilst also finding the poor people and making them watch *Jeremy Kyle* all day until their brains fall out. It may look like a meritocracy, but it clearly isn't.

Does talent behave like this in response to real-world experience ? I've no idea. The point is that we should be very, *very* careful indeed about interpreting any trends in real situations.

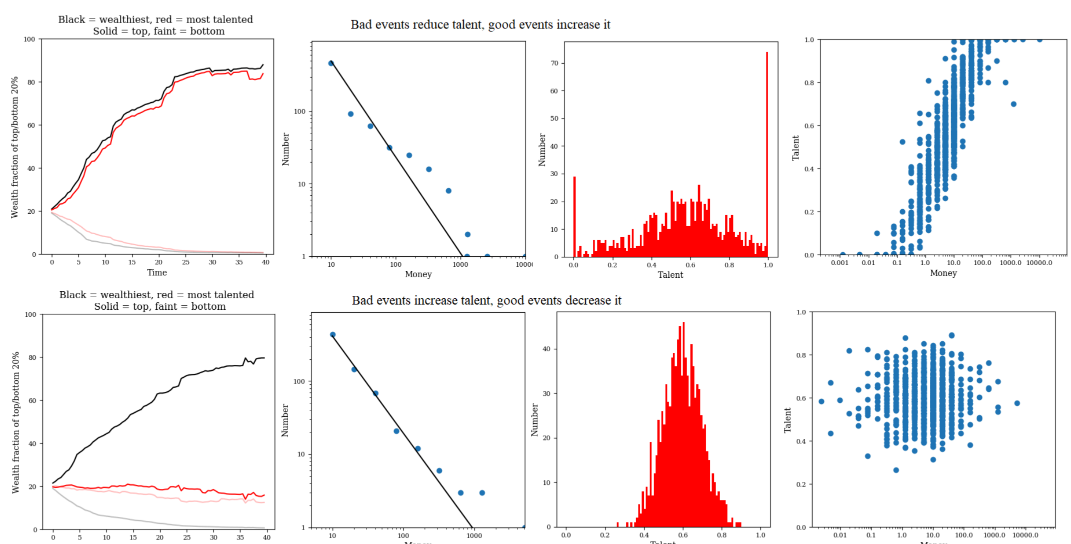

Time to go mobile

So far we've seen that the distribution of abilities and money can be connected in complex ways even under a simple model. It is possible to use a luck-dominant model to transform a Gaussian talent distribution into a power law distribution of wealth. It is also possible to use a talent-driven model to do the same thing : a genuine meritocracy may show the same behaviour. And if talent is allowed to vary, it is possible to create a perverted, false meritocracy where talent correlates with wealth but it's wealth which drives talent rather than the other way round. Interpreting this basic model is difficult because it is over-simplified and vague.

Previously we've tried varying the effects of talent on the encounters with events. The one remaining aspect is how agents seek out events. Recall the extreme case in which agents *experience* lucky events in proportion to their talent (and similarly avoid unlucky events). We previously saw that this causes a clear correlation between wealth and abilities, but there may be another way to produce the same behaviour : altering the distribution of agents/events and their movement direction. In reality there are clearly places where the probability of fortunate events is much higher : being in a research laboratory[8] with frequent discussions is much more likely to result in a breakthrough than sitting alone in a field. Of course you *might* make a breakthrough in a field, but it's not terribly likely[9]. So a talented, ambitious person isn't likely to be sat in a field for very long when they can see the bright lights of opportunity somewhere nearby.

Are talented people more ambitious ? I don't know. Again, people cannot be well-approximated by a two-parameter model. But since that's what we're going with, one thing we could try is allowing the agents to move in a direction determined by their talent. More talented agents could be more likely to move towards lucky events whereas the less talented could move at random (or away from lucky events, or even towards unlucky ones, *a la* the Darwin Awards). This *should* have similar effects to the talent-dominated models, but operates in a fundamentally different way. Rather than altering the nature of the events, which is possible for some situations but not for others, talent may be (perhaps more realistically) approximated by the skill to seek out opportunities and avoid events that will lead to misfortune.

In principle such a movement system could be implemented with the pseudo-random distributions used previously. We could grid the events and find the local maxima of lucky and unlucky events, and then have talented agents move towards lucky grid cells (or away from unlucky ones). But this would be quite complicated to actually do and likely very slow : for each agent we'd have to determine the nearest local maxima and find the angle to it, then check if there were unlucky maxima in the way… it'd all be just blaaargh.

The easiest way to test this is with a linear probability distribution that governs whether the events will be good or bad, e.g. a greater chance of finding good events on one side than the other. This make it trivial to have talented agents preferentially move in the direction of good events. The extreme case would be bimodal, in which there is a simple midpoint split : all events to the left are bad, all to the right are good (for example). Or there could be a linear variation, with there being some small chance of

[8] Or a library or a bank or a circus or whatever institution is relevant to the area of expertise.
[9] Unless of course the subject at hand concerns fields.

good/bad events even at the extremes, or perhaps only a smaller overlap region and zones where events are completely dominated by good/bad status.

It took a lot of systematic testing (i.e. trial and error and fine-tuning), but eventually I got something that works well. I tried using both linear and bimodal event distributions. The event movement, which is as in the fiducial run, is insufficient to significantly alter the basic initial distribution given the number of timesteps. At the end of the run, the linear distribution of events looks like this :

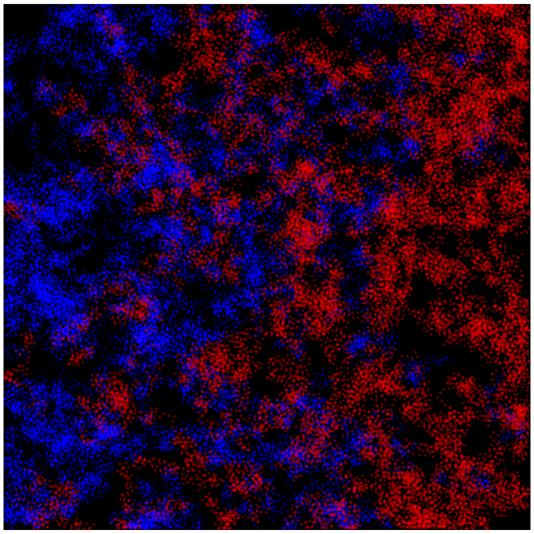

Where good events are in red and bad events are in blue. The closer an agent is to the right hand edge, the more opportunities for success they encounter. The bimodal split looks similar but with only a very small overlap region (caused by the random motions of the events) where both good and bad events can be found together.

I implemented the movement of the agents in the following way. First, their position shift is set to be 2.0 units in a random direction, which is exactly how the events move. But there is a chance proportional to the agent's talent that this will be overridden and instead they will move a given number of units directly left or right. Agents which have exactly the mean talent level have no chance of this - they always move at random. Agents with talent above/below some threshold from the mean are compelled to always move left or right, while the probability for agents of intermediate talent levels is interpolated linearly. Movement of the agents, like the events, is evaluated at the end of every timestep. As in the fiducial case, when agents intersect events they only influence the chance of exploiting lucky events and do not affect them in any other way.

I found that the most successful setup was to use the linear event distribution and a talent threshold for guaranteed left/right movement of 1σ deviation from the mean (if the threshold is set higher than this, then very few agents ever move directly), with a left/right movement rate of 5.0-10.0 units. I tried varying the event distribution, the movement speed and the talent threshold, which is a big parameter space to explore. So it's entirely possible - likely, even - that I haven't found the optimal case, but the end result is certainly good enough. I also tried control cases where the agents either don't move at all (i.e. the fiducial run) or only at random with no influence from their talent :

Linear event distribution, movement certain for T > 1 sigma

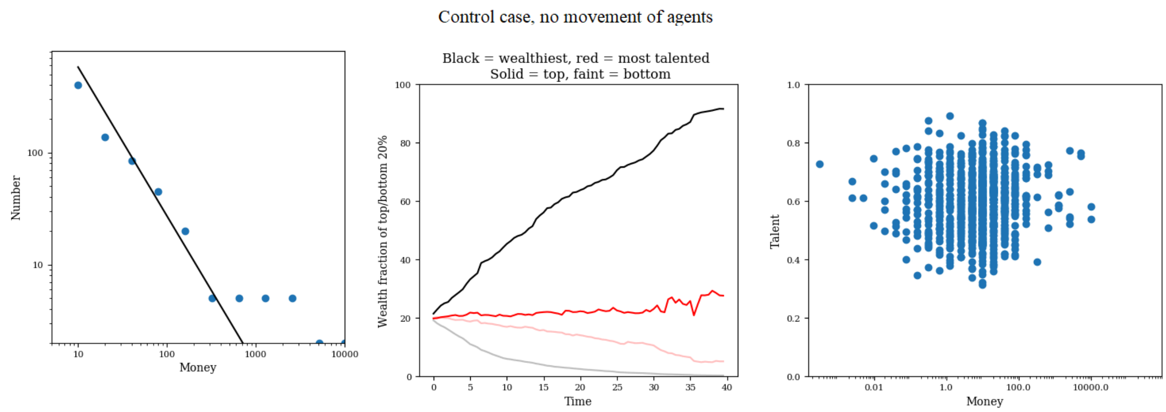

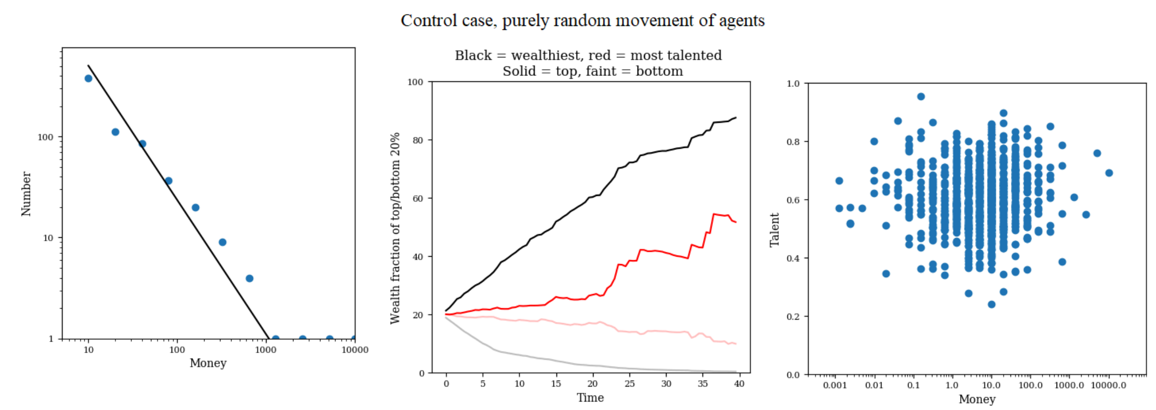

The purely random movement case is interesting since the 20% most talented agents become quite rich, yet 20% stupidest don't lose out very much. The Pareto plot is a bit stochastic, but the same behaviour happens in multiple runs (but not always - this might be pure random variation). There's no obvious wealth-talent correlation though. I'm not sure exactly what's going on here. If we allow the agent's talents to influence their movement then things get more fun :

Left/right agent movement 2.0 units

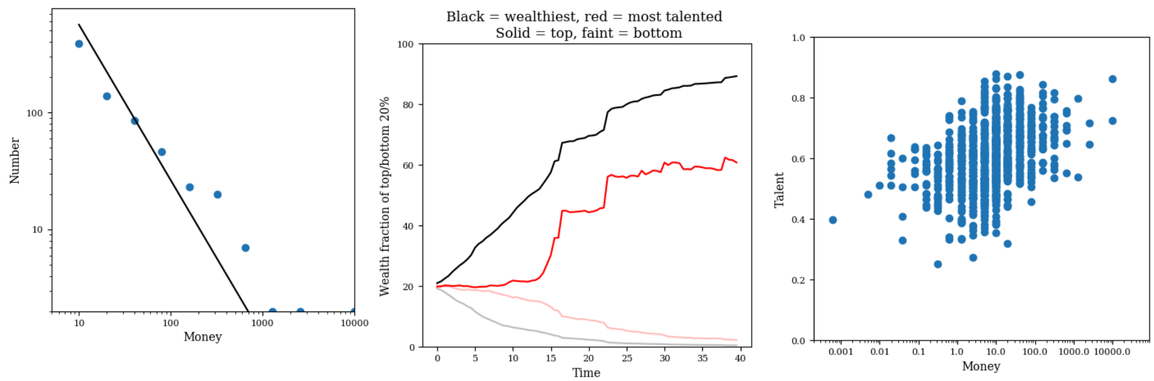

Left/right agent movement 5.0 units

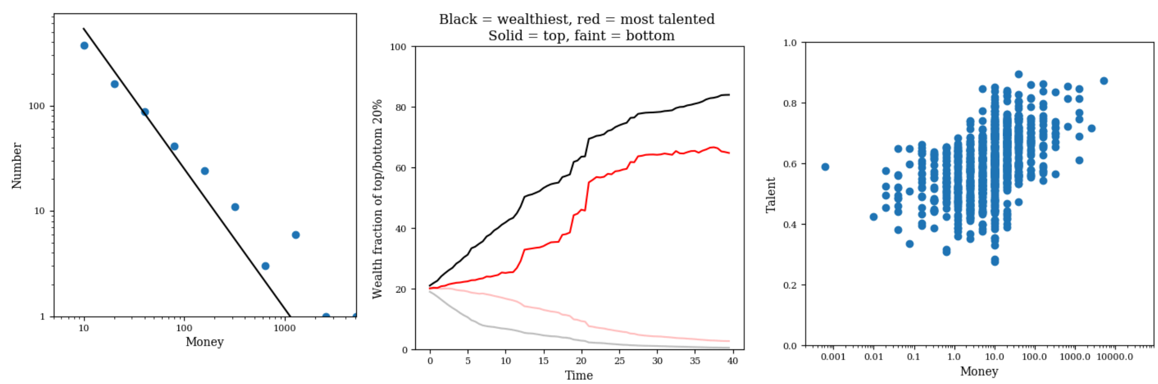

Left/right agent movement 10.0 units

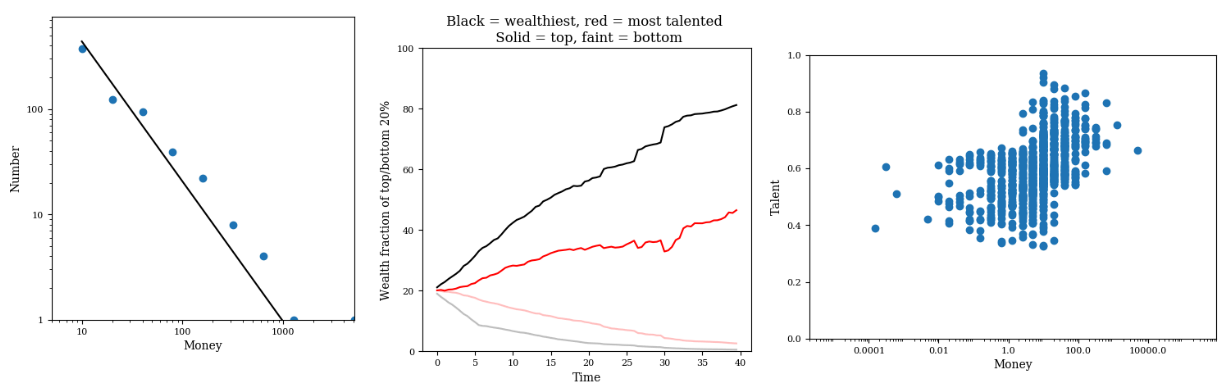

Left/right agent movement 15.0 units

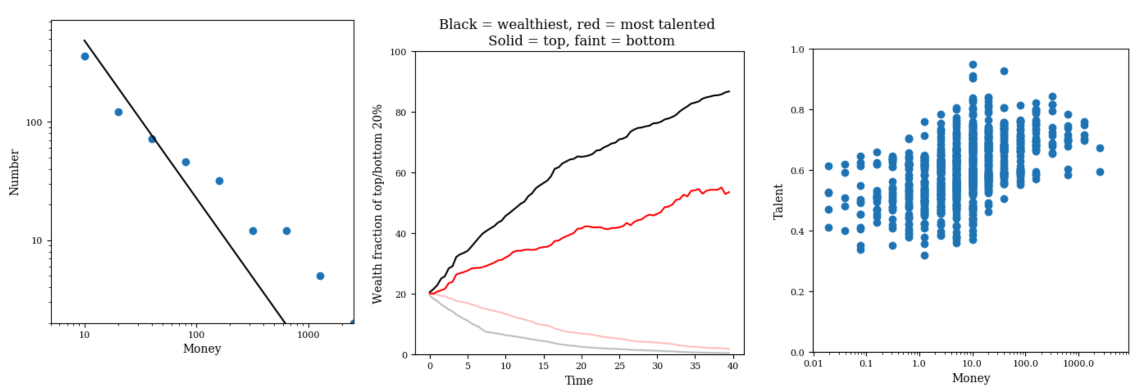

There is now a clear relation between wealth and talent, even if it doesn't always appear perfectly linear. The wealth distribution slope well-matches the fiducial case except when the movement becomes extreme. While not the perfect meritocracies we saw earlier, this works about equally well in comparison to the case of allowing talent to set the luck status of events encountered.

So this means there are at least two ways in which a genuine meritocracy can transform a Gaussian distribution of talent into a power law of wealth : talent can affect the type and effects of events encountered, or talent can drive people to preferentially seek out good or bad opportunities. Both happen to some degree in the real world. Of course, it's vital to remember that the models are only crude analogies to reality, not direct simulations of it. But these two methods are quite different, and independently undermine the conclusion (as does the false meritocracy scenario in which wealth affects talent) that luck is responsible for the difference in distributions. It could be, but there's no evidence to support that and other explanations work equally well.

(Somewhat surprisingly, using a position reset (as described earlier) does not really change any of the results for the different distributions : the wealth distributions simply flatline rather than increasing or decreasing. It may be that since the position reset is random, the agents which remain in the lucky areas effectively experience a simulation twice as long, so their increased wealth compensates for agents which end up in the slums)

All this may give the impression that I don't think luck is important. This is not the case, which I shall now elaborate at length.

Summary and Discussion

There are three areas this model has examined :

1) Determining the nature of real society based on observational data

PBR18 claim that while abilities follow a Gaussian distribution, wealth does not. The clear implication is that we do not live in a meritocracy, but we have seen that this is flawed for many reasons. First, there appears to be no evidence that abilities in general do follow a Gaussian distribution. It's quite possible that in some specific areas this is the case, but no evidence has been presented for this. Second, there is no reason to suppose that rewards are bestowed in direct proportion to ability[10], or that this has to be the case in order to have a meritocracy. Salaries based on ability but distributed as a power law would still constitute a merit-based system but give a very unequal distribution of wealth.

Even supposing that money was distributed in direct proportion to ability, however, it's by no means clear what a meritocracy would look like statistically. PBR18 used a model where wealth was driven almost entirely by luck. This gives a power law of wealth and no obvious correlation between talent and ability. But this work has reproduced their method and result and found that even in this case the situation is not as it appears. Even in PBR18's extreme model where talent plays only a single role - determining the chance an agent can exploit a lucky event - it still causes a correlation, it's just hard to see. The overall lack of correlation in the whole population occurs because the small number of agents (given their Gaussian ability distribution) severely limits the numbers of the most talented agents, and agents all experience different numbers of events. The larger number of mediocre agents means that it's much more probable that one of them will experience the greatest number of lucky events than a really talented agent. This makes it impossible to compare agents of a given talent as a population, because the numbers of high and low talent levels are so different. Yet the model makes it inevitable that the most talented *group* should do better than agents in groups of less ability; there should be no advantage to being mediocre, given the choice.

This is indeed the case, as can be revealed by two modifications : increasing the number of agents and/or changing the distribution of abilities from Gaussian to uniform. With those changes, we can plot wealth as a function of talent for agents who encounter similar numbers of events and find that there is indeed a correlation : the more talented generally out-perform the less talented. There is of course a strong scatter, because this model is undeniably dominated by luck. Being talented does not, in this case, bring with it automatic opportunities - those are down to chance, with talent only being able to exploit opportunities when they occur. But as a population, the most talented *always* win.

While it is undeniable that some real-life events are inevitably financially rewarding, and others, equally inevitably, financially damaging, it seems like a very extreme situation if talent does nothing at all except exploit potential rewards. This work has examined two other prospects : that talent alters the nature of the events encountered, and that it influences which events agents seek out.

[10] As mentioned, while there are some sectors where productive output might be a good, measurable proxy for ability (e.g. manual labour), there are many others where ability may simply be immeasurable.

Both of those scenarios are genuine meritocracies. The first allows an agent's talent level to determine the luck status (good or bad) of events it encounters (which still happens at random), while in the second scenario, I explored what happens if an agent's talent causes it to move towards areas of greater good or bad opportunities (whilst restricting it to only influencing the effects of good events when it encounters them). Both show a similar behaviour : a power law of wealth arises from a Gaussian distribution of abilities, but now there is a clear talent/wealth correlation, even if we use a small number of agents. Hence the statement that the different distributions of wealth and abilities is evidence of luck being the dominant factor is wrong. The different distributions is *consistent* with this scenario, but it is equally consistent with a meritocracy. There is no way to judge between them based on the distributions alone.

Even if we could measure abilities and wealth accurately, a correlation between the two does not necessarily indicate a causal connection. We know that such a connection exists in the above scenarios because we've engineered it. But we can get the same correlation in a very perverse way, by allowing talent to vary with experience. This is hardly an outlandish proposal. While some skills come naturally or very quickly, others are earned through practise. But in this model, it is of critical importance whether we allow good or bad events to influence talent, since the nature of the events correlates directly with wealth. This results in something which looks incredibly meritocratic, but in fact it's more like wealth (or equivalently luck) driving talent rather than the other way around. The analogy to the real world would be to find the rich people and given them highly specialised training whilst giving the poor people some low-quality brain surgery. Sure, there'd now be a correlation between talent and wealth, but it wouldn't be talent that had made people wealthy. It'd be the wealth that had allowed them (in effect) to buy an education. The message is that we have to be extremely careful about claiming a meritocracy (or otherwise !) based on observable trends : we have to be very sure which way the trend is driven.

2) The ethics of meritocracies - what do we really want to reward, and how much ?

We have shown that even in this model where luck plays an integral role (the encounters between agents and events is always affected by chance), broadly meritocratic societies can result. This raises the question of just how meritocratic we want a society to be. It's one thing to see the most talented people gaining the biggest share of the wealth, but another to see the least talented scraping the barrel. In the real world, the effects of wealth and talent are more subtle, and likely interrelated (as we saw when we allowed talent to vary). This model in which talent is a fixed quantity and should always be rewarded is clearly wrong, but this at least provokes a discussion as to what would be more realistic.

The obvious question is : what do we actually want to reward people for ? Clearly it's not just the mere ability to accomplish a task - they must actually do it as well[11]. So they need not only talent, but also motivation. Tasks which require a high ability should only be available to those with the necessary ability level or higher. But what if someone over-qualified takes a position ? Should we pay them based on their

[11] And yet, that can't be all there is to it. Most people can sustain an increased workload for a little while but beyond a certain point it becomes highly damaging. Yet others can persist at an increased rate for much longer. As far as I'm aware this is not because anyone has actively *chosen* how much work they can stand, but simply due to innate (at least partially genetic) factors. So do we want to incentivise working to extremes, thus rewarding a few individuals while actively harming the majority ? Surely this is not what payment is intended for. And yet at the same time, restricting those individuals who are able and willing to work longer feels oppressive and unnecessary.

prior experience (as employers often do) even though they're doing the same job that someone less qualified could do ? What if they're able to accomplish the same task more rapidly ? Performance-related pay *may be* valid in some cases (e.g. manual labour) but not in others (e.g. teaching, where there are elements that the employee cannot control). Even in manual labour, what one person finds exhausting another may perceive as no great burden. So should we reward the person who feels exhausted but gets the job done or the one who has put less effort into the task but accomplished the same or even more ? And what about how they spend their earnings - do we give the same amount to someone who spends their money, say, on campaigning to support fossil fuels or one who invests in green energy ? If the old saying, "they've got more money than sense" is literally true, then it follows we should give people varying amounts of money depending on how sensible they are… then again, the amount of money they have may affect their judgement. Should we really rely entirely on legislators to determine exactly what we can and can't do with our own things, or should this be considered at the point of employment as well ?

The situation is similar for the least talented. It is perfectly reasonable to suppose that if there are two *otherwise equal* candidates competing for the same job, the most able one should get it. The question is exactly how unequal a society is acceptable, and to what degree a lack of money is an incentive to work and at what point it becomes a crushing impediment to being able to develop necessary skills. The beatings surely should not continue until morale improves. But it would clearly be unfair to give two unequal people equal pay : for example, if we give them the same opportunities under the same circumstances, but one of them simply refuses to take them. The one who actually does the work and tries to improve obviously does deserve more than the one who refuses, even if the lazy citizen still deserves *something*. The root question, perhaps, is to what extent money itself creates opportunities, and consequently what the minimum standard of living should be to give everyone an equal opportunity for success[12].

There are limits to what society and individuals can do and expect : not everyone can (or wants to) become a virtuoso pianist, even with the best will in the world. We should not desire or expect equal outcomes from equal opportunities, because people are fundamentally different from each other. "Equal opportunities" might also not necessarily mean "identical situations", because a teaching method (or indeed raw subject matter) that works for some can be incomprehensible to others. And clearly there's a difference between someone who refuses to take an opportunity and one who tries but fails[13]. So if two people are doing the same job, but one person is unable even after expending great time and effort to perform as well as the other… why, exactly, should we more greatly reward the higher-performing individual ? Neither person *chose* their ability level. So why are we rewarding the genetic lottery, as well as other factors that people have no control over ?

People *do* make choices, they have agency. But it's clearly absurd to believe they can make *all* choices. A person born with no eyes cannot choose to grow eyes; not all skills can be learned by all people, even if they want to. The perception of merit always being awarded fairly risks the development

---

[12] I assume that there is an upper limit beyond which money won't help. E.g. a decent basic mathematics education requires "only" a good teacher and appropriate textbooks - if you can afford these, then adding in somewhat useful luxuries like hammocks with pillows stuffed with dodo feathers won't help, because student ability will be the limiting factor.

[13] And also one who tries, fails, but continues to do a crappy job when other things they could do much better are available (I'm looking at you, Mrs May).

of wealthism : the assumption that those who are poor deserve and/or want to be poor, that is *their fault*, not that of society. A belief that *only* individual agency and skill results in success absolves us from helping those without the opportunity to succeed; conversely, a belief that only luck and society affect the individual absolves us from taking responsibility for our poor choices and unjust actions.

None of this is to say that we can't at least sometimes make reasonable choices or that different wealth levels are inherently unfair. A meritocratic society is clearly an improvement on one in which wealth is bestowed according to creed, class or colour. But a pure meritocracy which only rewards success and throws the unsuccessful on the scrapheap - which in effect says that everyone is solely responsible for all of their own actions, as though every choice they made was perfectly informed and deliberate - is abhorrent. At the very least, such a system must be tempered by the knowledge that money itself provides opportunities to advance but also to corrupt. So what is it we want money to do, beyond providing the essentials of living ? How much freedom should we, as members of society, allow individuals to make their own choices with their own money ? Should we instead seek an arrangement where resources are distributed with the goal of greatest benefit in mind, or is such a goal beyond us ? Is the middle ground of a European-style blend of socialism and capitalism (or some other contemporary or hitherto tested system) as good as can be achieved ?

3) The philosophical nature of luck and talent

As has already been hinted at, talent and luck are not always independent. I.Q. scores are known to be at least partly due to genetics and also to correlate to some degree with overall success (both wealth and happiness, CITATIONS NEEDED). Similarly if someone is from a rich, caring[14] family, they will have more opportunities to develop their interests. That is not to say that exceptionally self-motivated individuals won't thrive in adverse conditions, but it is ludicrous to expect that everyone can do this just by shrugging off uncaring (or worse) parents. The child growing up in such conditions has little or no control over their circumstances - they therefore develop their own talents due in part to luck; their own abilities are also function of circumstance. This is true beyond childhood : employees supported in their development by their managers are similarly lucky. While as adults we can choose to some degree the environment we find ourselves in, there is always an uncontrollable element of chance. PBR18 provide a nice, convincing summary of why we often overestimate our own ability and downplay the importance of luck.

I have here restricted the discussion to financial wealth, which is much easier to measure than overall success. Yet clearly this isn't all people want out of life, even if they sometimes think it is. So the notion of a meritocracy must struggle with the further question of what kind of rewards people actually want and need. At least some studies suggest this is actually very modest in comparison to what they already have. While winning the lottery could be very lucky for some, for others it could bring nothing but misery. Chance is a relatively straightforward concept, but luck is highly subjective - as mentioned, a random discovery that would be fascinating and valuable to an astronomer would likely be meaningless and boring to a molecular biologist.

---

[14] CITATION NEEDED - I've heard that although there is a correlation between parent's wealth and child success, there's a stronger (and perhaps more causal) correlation with parental interest.

It's also sometimes important to *rely* on luck, in a sense. Any fool can repeat the same thing over and over on a different or larger sample, but to strike out on a genuinely new line of inquiry entails risk. At the same time, a stupid person could do the same thing : both luck and talent potentially resulting in the same discovery. The point about taking risks is that sometimes you have to fail, otherwise you were hardly taking a genuine risk at all. So luck requires failure, not just success. And the talent required to take risks therefore entails failure as well. How in the world do we decide how to reward failure ? Furthermore, the effects of events are hardly the all-or-nothing short-term occurrences modelled here : what seems - and may in fact actually be - wonderful when it happens can sometimes become extremely harmful on the longer term, and vice-versa. So whether we evaluate a chance event as lucky or unlucky is not only subjective but time-varying. Nor can we know for certain (in most cases) whether the alternative would have really been any better, adding still further to the subjectivity.

Beyond the interrelation between talent and luck is the question of how we evaluate talent. As discussed, most activities we reward have some associated quality factor. Art provides probably the most ambiguous example. What could be priceless to one can be perceived (quite literally in some cases) as rubbish to others. The financial value of art is dictated entirely by what people are willing to pay for it, regardless of the technical skill required for its creation. This surely comes closest to the facetious comment in the introduction that one sort of ability is the ability to make money. Even leaving aside the tricky problem of how we assess our own self-knowledge, how could we possibly judge someone's artistic ability ? No-one would seriously suggest we do this by seeing how rich they are (Van Gogh famously sold only two paintings in his life), yet that would seem to be precisely what happens.

Things are not that much better in the sciences. Again simplifying, suppose a scientist investigates an issue, forms a conclusion, and publishes a paper, all the while maintaining perfect objectivity. Do fellow scientists evaluate said research with perfect objectivity ? Of course not. They are influenced by current fashions, their own personal biases, the reputation and prominence of the researcher. Conscious of this, few would support the idea that citations are a foolproof guide to assessing research quality. Not only does the type of research vary with time, but just as with the arts, so does judgment of its quality - but not necessarily how the researcher themselves is perceived. Unlike artists, scientists are constrained by the available evidence. Provided they form a conclusion in agreement with that evidence at the time, they can still be regarded as a good scientist centuries later, even after their conclusions have been overturned.

These problems do not mean we cannot assess talent at all. Rather they caution that such assessments are prone to error, and all we can strive for is to make the best, most objective judgement possible - not to suppose that we have arrived at the true, exact value of something which really cannot be quantified at all.

A final point is that there are other effects at work which are completely different to talent and chance. Systematic effects mean that some career choices are far more likely to result in greater wealth than others, but people do not always even attempt to choose the path to greater wealth : they can, and do, knowingly select a job that brings greater enjoyment in itself than one which would be less enjoyable but provide more money. As we saw in the fiducial case, systematic effects mean that overall trends (or lack thereof) can be highly misleading - we must compare like-for-like; to compare the wealth distribution of footballers with scientists would disguise any trends within those respective fields (not to mention that comparing ability in these radically different skill sets would be meaningless). And the systematic effects cannot really be described as either chance or choice. From the individual's perspective, they *choose* (within limits of their abilities and opportunities) which career to take, but how much salary it pays is due to *chance*. From the view of society, salaries are only partially due to the choice of the employer - they are also due to the market forces, i.e. the choices of others. This is not chance, but neither is it choice - it is an emergent property, with no-one actively deciding to shape the economic system specifically to pay astronauts less than footballers or actors more than politicians. Our choices do matter, but their effects are complex, interwoven with the choices of others, and ultimately themselves at the mercy of random chance.

I think it would be foolish to try and answer some of the questions raised above. It is reasonable to suppose that if there are two otherwise equal candidates, the more able one should get the job. We might not be able to get a true measure of ability, but we ought to be able to get at least a useful measure of ability to do the assigned work. Beyond that, though, things rapidly become more complicated. While the simulations are instructive (and hopefully provocative), I think that they are a very long way from telling us anything about social policy as PBR18 imply. Rather I think they caution us that even establishing whether we live in a meritocracy or not is harder than we give credit for, and while clearly superior to caste-based systems, we ought not rush to assume that a meritocracy is what we want. Ability is not fixed : complex feedback loops and a host of other, quite separate moral factors necessitate that social policy has to be far more nuanced. And that is far beyond the scope of the current work.

What's next ? Possible extensions and notes to self to check existing conclusions

All the models investigated here have two main parameters for the agents : talent and money (position notwithstanding, that's more of a zeroth parameter). This has been investigated pretty thoroughly, I think, and seems like a sensible cutoff. Any future extensions to the project could consider extending this to add more parameters or change the fundamental basis of the model in other ways. The complexities of even this absurdly over-simplified two parameter model ought to give anyone claiming they understand how real society works serious pause for thought…
Items labelled as "to check" are ones that should be investigated before writing this up more formally. Unlabelled items are ambiguous.

- To check : for the case of varying talent levels, compare the wealth evolution of those who were most talented at the start with those at the end.
- To check : reset cases for the other event distributions, with and without agent movement.
- To check : movement with linear event distributions and talent affecting unlucky events.
- To check : for the case of resetting the agent positions halfway through, do the Pareto plot for the wealthiest at the midpoint for the whole duration (see if those were started off as the richest for the second round remained there - they almost certainly do, but this should be checked). CONFIRMED, THEY DO.
- To check : do the randomised positions halfway through for the case of different distributions of events. This should have a much stronger effect, i.e. more chance of a wealthy individual suddenly being surrounded by bad luck (especially for the bimodal case). NOT REALLY. WHICH IS A BIT SURPRISING.
- To check : redo earlier scenarios with the standard wealth cap and Pareto plot that includes the most/least talented and wealthy
- Future work : add additional agent parameters : ambition, wisdom, etc. In reality it's not those of sheer ability who succeed, but those who actually do the work (as well as those who are just lucky).
- Future work : allow multiple event types, some which are influenced by talent and some which cannot be.
- Future work : investigate how wealth varies with geographic area; more complex event position distributions… haven't tried varying event density.
- Future work : you can't get anywhere without wealth. Mobility costs money. Money-making resources themselves cost money. You do need some amount of talent to select sensible investments but you also need a load of cash.
- This model only considers a population of a fixed uniform age. In reality the measured wealth distribution doesn't do this. Run multiple generations at once and estimate the distribution from the whole population rather than from one generation.
- Can we make a system where even if we start with an initially inverted distribution of wealth (such that the most talented start off very poor and vice-versa) the most talented people eventually become the richest ?