Bottom line

I think that log(wealth) is a good first pass utility function for altruistic investors.

An investor with log utility should leverage significantly on average, probably 2-3x if they invest only in stocks and significantly more if they invest in a lower-risk diversified portfolio.

I think choosing optimal leverage can increase log utility by >2%/year for a stock-only investor and >3%/year for a diversified investor, though I don't actually run those calculations here except for stocks at a single "representative" time.

There are a lot of complications to think about, but I think the basic story is fairly compelling and not very speculative. I've been doing this for the last few years.

How to think about risk

For relatively liquid assets, risk-tolerance can be basically quantified by a single parameter "how much you like risk" which determines how much leverage you should use in a direct way. This isn't obvious at all, I get it into it in an appendix.

How risk averse should we be?

I think people should default to logarithmic utility. This corresponds to a simple rule, which I'll justify in the appendix:

> Leverage = (expected excess returns beyond the risk free rate) / (variance).

There are two kinds of reasons to use this rule:

- It seems kind of plausible on its face, at least if I'm a really big donor. Once I'm spending $10 billion on a problem, it seems plausible that it's about 10x harder to make further progress than if I was only spending $1 billion on that problem. When thinking about causes like GiveDirectly I don't even think returns diminish that quickly. But logarithmic utility already corresponds to a very leveraged position, so I think it's reasonable to at least start there.
- Over the long term, using logarithmic utility maximizes your median wealth. In fact, over a long enough time period agents with logarithmic utility almost surely end up with more money than agents who are more risk-averse *or* agents who are less risk-averse. More risk-averse agents just grow more slowly, while less risk-averse agents have their wealth converge to 0 almost surely. So I think logarithmic utility is a pretty natural default even if in some sense you are agnostic about how fast returns diminish.

I’m not getting into the correlation between market returns and the availability of philanthropic opportunities. I think that this doesn’t change the picture very much for big donors but is extremely important for small donors. (It’s not obvious that it doesn’t change the picture, but that’s a longer discussion I don’t want to get into here.)

What does this say to do?

First build your optimal portfolio.[1] Then estimate the expected returns (after subtracting the risk free rate) and the variance. Take the ratio. That’s your leverage.

I go through the math in an appendix.

In the rest of this post I’m going to help ballpark this by talking about stock market returns and variance. In reality you will do better by having a more diversified portfolio with more leverage. But I think this is still very helpful for thinking about the general ballpark of how much risk you should be taking on if you have logarithmic utility.

Risk-free rate

This is easy to look up, just see what interest you are paying on your margin loans (or getting paid on your cash reserves, if your leverage is <1).

I’m going to assume that the risk-free rate is a constant at which you can both borrow and lend. In reality there is probably some gap there no matter how efficient you are. It doesn’t really change the picture though. Use whatever rate reflects the marginal value of $ in your bank account.

Right now the risk-free rate is around 0%. It can be hard to borrow at such low rates, though you can still get embedded interest rates this good from futures and options. When risk-free rates are higher it’s a bit easier to actually borrow at them and you don’t need to jump through hoops.

Before the crash risk-free rates were around 1.5%/year.

Estimating returns

I’m in favor of using pretty simple heuristics for estimating market returns. I don’t think there are market forces that tend to make these simple estimates systematically wrong when applied to the market as a whole (though there are if applied to individual sectors or companies).

---

[1] I like the market portfolio + adjustments for my contrarian beliefs. Use whatever you like though, everything will be the same but with a more aggressive expected return estimate. In general I think if you are making contrarian bets then you should still *also* get market exposure and at the end of the day should usually end up even more levered than normal (even if the contrarian bets are quite risky).

I think that simple heuristics suggest that current[2] returns 7%/year returns for the global stock market right now (nominal CAGR). I describe this estimate in an appendix. You would estimate other markets differently. I think stocks and bonds are most important, and bonds are somewhat more straightforward to estimate.

(This number was more like 6.5% before the crash. And the estimate I'm using relies on a few assumptions---~0% expected earnings growth, ~0% expected multiple change---that are unlikely to hold during the crash but which point in opposite directions.)

Estimating variance

For most stocks and bonds, people sell options. Option prices tell us something about the distribution over values that the asset might have in the future, and from those we can infer a variance. VIX summarizes this for the S&P500 (VIX is the annualized standard deviation, so you need to square it to get the variance), though you can also compute it for most anything.

I don't trust these numbers that much---they reflect hedging value of options *etc* so I think tend to be a little bit conservative.[3] But I still end up using them to decide how much leverage to take on. If you disagree with the numbers a lot, you are probably better off trading in options directly rather than using substantially more or less leverage.

Right now VIX is at a historical high of 57.24, which corresponds to variance of 32.7%/year. Long-term averages are more like 15, which correspond to a variance of 2.25%/year.

CAGR vs arithmetic mean

In the formula for leverage we wanted to know the expected returns. But when I say "this asset neither goes up or down over the long term" I mean that if you take the *geometric* mean you get roughly ~0%. If you went up or down 10% a day every day, then within a year you would almost certainly have fallen by more than 75%. So a quantity that is stable over the long term should have a positive arithmetic mean.

The size of this gap is roughly half the variance. So the CAGR vs arithmetic mean adjustment causes me to increase my leverage by 0.5. This isn't a huge deal most of the time, but during this crash it's really important. Because it's not normally that important, I haven't thought about it that much.

---

[2] April 1, 2020---this doc is written at this moment in time, but I'm going to try to make some of the calculations from a more representative time to show more what normal numbers look like. In retrospect I should have just done all the calculations for a more representative time.

[3] For example, option prices imply that all assets have roughly zero excess returns---this is necessary or there'd be an arbitrage. But I think the variance estimates are closer to correct.

Putting it all together

I'll sketch what this gives during "normal" times (like 6 months ago) with an estimated CAGR of 6.5%, risk-free rate of 1.5%, and implied variance of 2.25%/year. In a later section I'll talk about right now, which is a bit of a weird time.

Here are the expected log returns for different amounts of leverage, calculated in [this spreadsheet]:

- 0x: 1.5%
- 0.5x: 4.28%
- 1x: 6.5%
- 2x: 9.25%
- 2.72x: 9.84% (optimal)
- 3x: 9.75%
- 4x: 8%

You get most of the benefit from getting your leverage even roughly correct---you get 3/4 of the gains by moving 1/2 of the way from where you currently are to the optimal value of leverage.

Taxes and leverage

The interaction between taxes and leverage is kind of complicated and depends on details of your situation. In general, I think that considering taxes tends to increase your desired level of leverage. If you also donate, then the effect becomes even stronger.

These are strange times

Right now, that calculus suggests a CAGR of 7% and a variance of 32.7%. This is really a deeply and qualitatively different situation. Without the CAGR adjustment you'd end up with leverage of about 31%, i.e. you'd put 31% of your money into stocks and the rest in short-term debt. With the CAGR adjustment, you'd have leverage of 81%.

This roughly corresponds to the belief that the market is a random walk---if current conditions prevailed for a long time, it's expecting prices to trend upwards by only 7%/year, while there are swings of almost double-or-alf every month.

The CAGR adjustment is pointing out that if the market isn't drifting up or down, then a portfolio with leverage between 0 and 1 will actually tend to drift up: every time the market goes up you'll sell a little, when it goes back down you buy a little, and on average you rise.

I think that’s plausible and it’s roughly how I’m investing right now, but realistically I think the current conditions are weird enough that you might want to just be opinionated about what’s going to happen.

My default expectation is that volatility will continue falling and the strategy in this doc will recommend levering up gradually over the next few months. I don’t think it will be that long before recommended leverage is back to 200% (roughly would require VIX back down to 25), and I think this will likely be a relatively important time to get back into the market with leverage.

Appendices

Appendix: Dynamic leveraging, intertemporal betting against beta

I recommended using option prices to estimate volatility. Option prices change over time.

There are two ways to do this strategy:

- Pick some long horizon, estimate volatility, and stick with that volatility estimate. This sort of bakes in the assumption that short-term higher volatility comes with higher returns and everything works out to be fair, which is what would happen in EMH-world.
- Constantly update your volatility estimates and change your leverage when VIX changes. This is what I do, but it's an anti-EMH bet.

I think it's reasonable to do things either way. These two policies recommend similar average levels of leverage, though the dynamic one is slightly more aggressive.

If all investors levered up when volatility is low and down when it's high, then market prices would tank when volatility goes up and then predictably rise when volatility goes back down, and this strategy would be systematically buying high and selling low.

The CAGR correction does say that expected returns are higher when volatility is high, but it's not that big an effect---right now that's saying that we expect 16% excess annualized returns as compensation for the variance.

In EMH-land the excess returns would be large enough to fairly compensate for the increased volatility, such that no one wants to buy or sell. Interestingly, this holds even if investors have different risk tolerances.

There are a few reasons I think the world probably isn't that efficient, such that I think the dynamic approach is probably better and the CAGR correction is in the right ballpark.

- Betting against beta is the regularity that risky assets tend to systematically outperform leveraged low-risk assets. I think this is one of the stronger market regularities, and it also has a pretty good story---there are a ton of reasons that investors can end up leverage-constrained or leverage-discouraged (mostly because it's a more visible signal of risk-taking than high risk assets), so people who want risk end up bidding up the price of high-volatility assets. Dynamic leveraging is basically the intertemporal version of this, betting that the excess returns during high-volatility periods are moderate and don't fairly compensate for the risk. When I look concretely at the world, this basically rings true to me.

- Empirically, high volatility times haven't had significantly better returns. The average has actually been negative, though the standard error is about as large as the effect we're looking for so it's mostly noise and mostly helps for ruling out even bigger impacts.[4] There is enough noise that we can't rule out high volatility times doing slightly better, but I think we can pretty decisively rule out the hypothesis that there are enough excess returns to fairly compensate the extra risk. I would like to see a more careful version of this analysis, and I'm not that confident in these conclusions.
    - This was my preliminary conclusion from a backtest in 2017.
    - It was true while I've been invested from 2017-2020, leading to me moderately outperforming (I think >1%/year though haven't been great about tracking the impact of my various bets separately).
    - This policy did incredibly well during the coronavirus crash (and gets a lot of its backtested oomph outperformance from similar crashes), leading to something like 40% excess returns.[5]

---

[4] It's a little bit subtle how to think about error bars in this analysis. I think the biggest issue is that the situation today may be different from the situation in the past. But I think this is a case where I don't really expect the market to "fix" the inefficiency, as described in the next bullet point.

[5] On February 21, the Friday before the crash, this strategy was already down to 2/3 of its normal leverage. On February 23, after the market fell 3%, this strategy sold more than half of its stock. On average, it held <20% of its normal exposure during the crash, and is now gradually inching its way back up. I also took a smaller inside view bet that coronavirus would be bad, but it was smaller than the outperformance of dynamic leveraging.
(I wrote that on April 1, by April 15 the market has recovered and some of the gains are gone, though it's still way up and roughly as important as my inside view bets. Wouldn't really want to update my reckoning for dynamic leverage until volatility is back to normal levels.)

Appendix: margin limits, margin rates, brokers

Most brokers effectively don't offer margin accounts by having totally absurd margin rates.

Interactive brokers offers margin rates that are roughly the risk-free rate + 0.3%, and has a margin limit of about 10x if you opt for a portfolio margin account (i.e. you can borrow about 900% of your net worth, though it depends directly on the risk you are taking on). Those limits are high enough that an investor with logarithmic returns will basically never run into them.

I've generally found interactive brokers really nice to work with. I think that most large investors probably have better solutions than this.

You can also effectively borrow by using options and futures. I haven't done this, it can probably get you slightly better rates (significantly better when interest rates fall below 0.5% because even IB doesn't offer good margin rates in very low-rate environments). I've used explicit margin because it seems easier. I've also recently learned that margin loans (but not embedded leverage) interact favorably with deductions for charitable contributions.

Appendix: logarithmic utility --> (returns) / (variance) leverage

Logarithmic utility is particularly easy to think about, because the log of my returns over a year is just the sum of the log of my returns over individual months. That basically lets me think about every month independently (and similarly every day, or year, independently). So I don't need to think about exactly when I'm spending my money or predict how much money I'll have then. Obviously things would get more complicated if e.g. I knew that the first billion had great returns and then returns diminished quickly after that, but I think that given a lot of uncertainty we should be more skeptical of that kind of model and move more towards something like log returns.

So suppose that I can borrow or lend at the risk free rate r, and I have an investment that returns X over the next year. Write m for expectation of X v for the expected square of X[6].

(Things would be identical if we chose a shorter period, and it's kind of important that we rebalance every time the market moves a ton otherwise e.g. we can get totally wiped out.)

I'm going to talk about everything in units where 1 = all of my current money, to make life simple. That's another way logarithmic returns are simple---the optimal strategy doesn't depend on how much money you start with, we can rescale things freely.

---

[6] Expected returns squared is basically equal to the variance as long as I think about a short enough time period: expected squared returns are equal to variance plus (expected returns) squared. If I cut my interval in half, then the expected returns and variance both get cut by 2. So the (expected returns) squared term gets cut by 4, and if I keep shrinking the interval it eventually becomes unimportant.

If I leverage this portfolio up by a factor of k, then at the end of the year my money will be:

$$1 + r + kX - kr$$

The 1 is the money I currently have. The first r is the interest I earn on that money. The kX are the investment returns I get. The final kr are the interest I pay to finance my investment (or the interest I forego by loaning out my money).

So now my utility is: $\log(1+r + k(X - r))$.

The first r out front can be basically ignored if we are considering a short enough period, it just becomes an additive constant. So then we are left with $\log(1 + k (X - r))$. As described in the next appendix, we can approximate this to second order:

$$\log(1 + k(X - r)) \sim k (X - r) - k^2 (X - r)^2 / 2 = k m - k^2 v / 2$$

The value of k that optimizes this is $k = (m - r) / v$.

Appendix: math about risk tolerance

How much leverage to take depends directly on your preferences about money. If gaining \$1 is exactly as good as losing \$1 is bad, then you shouldn't care about risk and should just maximize expected returns.

It turns out we can basically think about risk-tolerance in terms of a single parameter: how much would you pay to avoid a gamble that has a 50% probability of gaining \$1 and a 50% probability of losing \$1? This is basically about diminishing returns: how much worse is dollar N+1 than dollar N?

If I know that, then I can basically work out how much I'd be willing to pay to avoid other gambles. For example, if I took four \$1 gambles, I'll probably end up gaining or losing about \$2. So I should be willing to pay about four times as much to avoid a \$2 gamble as a \$1 gamble. In general, the badness of a gamble is *quadratic* in its size. This is really important.

(For a sense of perspective: I think that if you have about \$100M, you should be willing to pay about \$5,000 to avoid a gamble that will cause you to gain or lose \$1M, or \$500,000 to avoid a gamble that would cause you to gain or lose \$10M.)

Formally, I have some utility function U(dollars) that tells me how happy I am if I end up with a given number of dollars. The slope of this function is how much I value a dollar. The curvature tells me how bad a gamble is---how much higher is U(n) than the average of U(n-1) + U(n+1).

Locally I can expand this utility function as U(current wealth + d) = x_0 + x_1 * d + x_2 * d^2 + x_3 * d^3 + ...

If I'm willing to rebalance my portfolio every time the market moves by more than 10% or so, then it turns out I can basically ignore the third order terms, so I only need to think about x0 + x_1 * d + x_2 * d^2.[7]

x_0 and x_1 don't really affect anything about my investment decisions, that's just an arbitrary shift and scale of my utility function. The only thing that matters is x_2 / x_1. Once I know that, then I can set my optimal portfolio.

It turns out that things are even easier than this: I can compute the Sharpe-ratio-maximizing portfolio given my beliefs, and then the ratio (x_2 / x_1) will just tell me how much to leverage up that portfolio via a pretty simple formula (leverage = sharpe ratio * (x_1 / x_2)). So I can basically decouple the decision of how to invest from how to leverage (unless my portfolio has a ton of illiquid assets that I can't rebalance easily).

Appendix: Estimating returns

I think returns of stock can be pretty well approximated as:

- Earnings yield: how much do the companies actually make this year
- Earnings growth: how much more will these companies make next year, without new investment?
- Multiple changes: how will the P/E ratio for this stock change over the next year?

The first factor is easy to ballpark. E.g. for the S&P 500 it's currently about 5.4%/year, the inverse of the P/E ratio. You tend to get larger numbers in international markets, but I usually use US markets to estimate returns (and assume that prices are fair across markets) because I have considerably more faith in US corporate governance and accounting.

The second factor is hard to estimate. On average across all companies at all times it's been something like 0-1%/year above inflation. I think it's more like 0% or slightly negative during good times and becomes significantly positive during recessions depending on the exact nature. I think using inflation is a fine proxy. Current ten-year break-even inflation rate is 0.87% (inferred from TIPS prices, basically pure bets on in the inflation rate), though I think this is systematically distorted during crashes and I'd use more like 1.5% right now. So that brings us to 7%.

---

[7] It's a little bit weird that you can ignore third order terms but not second order terms. This is because of a pretty deep fact about how stochastic processes work. It's approximately true as long as (i) there is some bound (like 10%) on how much the market can move instantaneously, (ii) regardless of what's happened up until this point, the expected return of the market over a short time period is approximately zero (e.g. you can't predictably increase your money by more than 10% over the course of a single day).

P/E multiple changes: on average over very long time horizons these aren't big. In 1920 multiples were around 10 and 100 years later they are around 18.63. That wasn't a secular trend either, mostly just noise where 100 years ago happened to be unusually low. But even for that extreme measure you are still only at 0.6%/year. I think that the expected impact of these changes is pretty small unless you want to be opinionated about changing risk premia or interest rates, so I usually call it 0.

I actually expect the *geometric* mean of returns (called CAGR) to be around 7%---if you bought and held for a long time, and the market looked roughly like this, I'd expect your money to go up by around 7%/year on average. That means the *arithmetic* mean is somewhat higher. We'll get to this below.