INVESTMENTS & WEALTH MONITOR

A reprinted article from May/June 2020

THE NOVEL CORONAVIRUS PANDEMIC

A Quantitative View

By Dan diBartolomeo

INVESTMENTS & WEALTH INSTITUTE®

THE NOVEL CORONAVIRUS PANDEMIC

A Quantitative View

By Dan diBartolomeo

Editor's Note: This article is based on a combination of two articles, both of which are posted on the author's website: "The Corona Virus Emergency: A Quantitative View" (March 20, 2020), https://www.northinfo.com/Documents/937.pdf and "Update on the Coronavirus Pandemic: The Status of Credit Markets" (April 16, 2020), https://www.northinfo.com/Documents/941.pdf. Information in this article is updated as of May 7, 2020.

The coronavirus pandemic represents an unprecedented challenge to all investors. As of this date, the global situation remains deeply uncertain and subject to change in conditions by the day. In the current atmosphere of an ongoing crisis, it is critical that we help investors have a logical rather than an emotional basis for their actions. The first purpose of this article is to examine the likely outcomes from the current pandemic through a simple first-order model. The second purpose of the article is to discuss the related implications for investors.

This model is not intended to compete with forecasts arising from public health authorities; its purpose is to frame the problem in a way that is transparent to financial professionals so the output can be integrated easily into their investment decision-making. Many parameters of the pandemic model remain highly uncertain, and the available body of statistical data continues to change from hour to hour.[1]

The worst-case scenarios for global coronavirus mortality are horrific. However, in the course of the 20th century, humanity survived several incidents that were even more extreme (World War I, 1918 Flu Pandemic, World War II). Our estimate of the worst case was that the entire world ignores the pandemic (obviously already untrue) in which everyone in the world becomes infected with the coronavirus, leading to mass mortality of around 60 million.

The estimate of 60 million deaths reflects the fact that the 3.4-percent mortality rate (a March 3, 2020, estimate of the World Health Organization) is highly skewed to an elder population that is large in the countries impacted as of that date. For example, the United States has about 4.3 percent of the world population but about 12 percent of the world population older than age 65. When you break the observed mortality rate into the "under 60" and "over 60" subsets, you capture the fact that the percentage of persons older than 60 in the developing world is quite small. There are approximately 815 million people 65 and older in the world or about 10.5 percent of the total population. If the daily growth rate of new infections within the 65 and older age group alone can be brought down to a low level (e.g., <2 percent), the total mortality from coronavirus reduces to a very low level. Later we will consider several different ways in which the 3.4-percent mortality rate might be influenced to be either very much lower (a small fraction of 3.4 percent) or as high as 14 percent.

Although terrifying, the 60-million casualty estimate is about 0.8 percent of the current world population. Although casualties of war are hard to tally with precision, that figure of 0.8 percent is around the higher end of the estimated range for each of the eight-plus years of the World Wars. When we add in the massive number of deaths (median estimate 35 million) from the 1918 Flu Pandemic, the maximum likelihood estimate for abnormal mortality for that year was about 1.9 percent of the global population at that time.

Let's try to put the coronavirus emergency into this context. As of mid-March 2020, about 10,000 people had presumptively died from the coronavirus since the first apparent case on November 17, 2019, in Wuhan, China. Over that same span of about four months, about 19 million people had died from routine causes (old age, disease, traffic accidents, and armed conflict) unrelated to the coronavirus, out of the current global population of about 7.7 billion. By early May 2020, the death toll stood at about 250,000, but we estimate the total number of deaths globally at about 26 million over the same period. The current impact of the virus on global mortality has been small at about 1 percent of all deaths, though the figure has approached one in two deaths in a few cities at the respective local peaks of the crisis.

To summarize the developments subsequent to mid-March, we can simply say that the best-case scenarios have gotten much worse, but the worst-case scenarios have gotten much better. Put simply, the process of slowing the spread of infections in Western Europe and the United States has taken much longer than in the reported data on countries earliest to be affected, China and South Korea.

On the positive side, almost all the countries where the pandemic has infected large numbers of people to date are now reporting declining numbers of new infections and stable or declining numbers of deaths. Because the typical period of the illness is a couple of weeks, it is to be expected that trends in mortality will lag infections (both known and unknown) by a similar interval. In some countries, sporadic spikes in new infections simply reflect increased rates of testing that results in a counterintuitive but positive sign. Finding persons who are infected implies there are fewer people "infected but unaware" (who are most likely to spread the infection). A more detailed discussion of testing rates across a range of countries appears below.

Even stable numbers of new cases imply an end to an exponential growth process that could have created tens of millions of deaths if left unchecked. For example, a U.S. Centers for Disease Control (CDC) estimate as of late February of the "worst case" for U.S. mortality was 1.7 million deaths.[2] This huge number was later increased to 2.2 million in a separate analysis that assumed that only 50 percent of the American public would abide by social-distancing orders. Luckily, the rate of U.S. public compliance appears to be well over 90 percent nationally.

At the start of May, the current upper confidence bound on U.S. mortality from the widely cited University of Washington IHME models[3] was around 115,000, less than one-tenth of the CDC estimate of two months ago. The central tendency of the U.S. mortality estimate from the IHME model was about 72,000 with minimal mortality after June 1. Although still horrific in human terms, it is of the same order as the 61,000 U.S. deaths arising from regular seasonal flu during winter 2017–2018, despite the existence of a vaccine. In May, many U.S. states joined some European and Asian countries in loosening lockdown restrictions. Subsequently, the IHME model was updated to include an expected increase in infections associated with greater societal mobility. The upper bound for expected U.S. mortality from this model has about doubled to 240,000, which is still less than 15 percent of the CDC worst-case scenario of February.

From a human perspective, the vast majority of risk is how great the loss of life would be if the virus is left completely unchecked and how soon. At a geometric growth rate of 8 percent per day (see discussion below for description and distinction from related statistical parameters), the entire world population would be exposed in approximately four and a half months. If the growth rate of exposure is reduced to 1 percent per day, fewer than one in a 1,000 people would be exposed in the next year and roughly one in 30 people would be exposed in the next two years, which is the widely suggested estimate of time to availability of a vaccine.

As of mid-March, the decline of the S&P 500 since the recognition of the pandemic implied that equity investors believed at that time that the expected duration of the crisis was on the order of seven months. By mid-April, the expectation of the duration of the pandemic had declined to five months, suggesting an expected return to relative normalcy in September 2020.

A FIRST-ORDER MODEL FOR THE PANDEMIC

I present a simple model for tracking the pandemic.[4] There is a key subtlety that must be understood about the model and why it is different than the range of pandemic statistics being widely reported. The real issue isn't the number of cases we know about, it's about getting a reasonable estimate of the number of cases we don't know about, because those are the people who are continuing to widely spread the infection.

As countries ramp up measures against the spread of the virus, the true but unobservable growth rate of the infection logically cannot be going up. This means that if the reported growth rates are going up, it has to be because we're discovering a larger portion of the existing pool of previously unknown cases. The faster that countries discover the previously unknown infections, the faster that the spread eventually will subside, because persons who are known to be infected can be more effectively isolated, greatly reducing the spread of infection. This leads to the somewhat counterintuitive result that rapid growth in the number of reported infections may be a positive development in that the number of existent but unknown infections is apt to be declining.

The model uses three key input variables: the mortality rate of those infected (M), the rate of spread of the infection per day (E), and the discovery of infected persons through testing, which impacts the number of patients actively in treatment (A). A fourth parameter, the capacity of the healthcare system, is presented as a condition of the extreme upper boundary on mortality. For any selected values of the parameters, the potential for the virus to spread from infected to uninfected persons can be estimated along with a forecast for cumulative mortality.

Here is a set of stylized facts to start, which should motivate the design of the model. It is now believed that the first COVID-19 case appeared in Wuhan on November 17, 2019. By January 22, 2020, 845 cases had been identified and an unknown number of additional people had been infected and were unaware they were infected. For those first 66 days, this is an exponential growth rate of about 10.7 percent per day (noted as E in the model). Given that many infections probably went undetected at the time, the true rate of spread might have been much higher in those early days (maybe 15–20 percent per day). There was also a spike in reported cases between January 22 and January 31, when widespread testing went into effect, with a daily growth rate of about 35 percent as previously undetected infections became known. Chinese health authorities also changed the criteria for presuming someone was infected in mid-February, which also caused a one-time spike in the apparent growth of cases. It should be noted that Wuhan is a city of 11 million people with an area of 3,820 square miles. This works out to a population density of about 3,000 persons per square mile. At the higher end of the density scale are older cities like New York (28,000 per square mile) and London (15,000 per square mile). By contrast, the population density of the world is about 317 persons per square mile of habitable land.

Please note that the current daily exponential growth rate "E" is not the same mathematical quantity as the exponential growth rate of an infection across patients. This value is widely referred to as the parameter "R" in medical studies. The two concepts are similar and are a frequent source of confusion. The R parameter refers to how many new infections might be caused by interaction of an infected person and uninfected persons. For example, if R is 2, then one infected person is expected to infect two others. Those two infected persons would be expected to infect two each, or four. So, in two of the transmissions from a single infected person the count has grown to seven (1 + 2 + 4). Although the R-value for the coronavirus has not been well-established, early indications were closer to R for common flu (R=1.2) than for some other illnesses. Later studies suggested R estimates in the range of 2 to 2.5, which is comparable to Ebola, but still far lower than for measles (in excess of 15).

To get a useful estimate of the likely extent of the pandemic, what we really care about is how many people are continuing to spread the virus. This consists of two groups, those that don't know they are infected and those that do know but continue to interact with others. The mathematical model[5] is designed to infer the number of "infected but unaware" persons. We define a particular time interval of interest (e.g., two weeks) and then estimate the number of deaths expected in that period based on the information available at the start of the interval (number of known infections) and the assumed mortality rate. If the number of deaths during the interval is larger than was expected, the additional mortality is presumed to be the result of infections that had not been identified by the interval start date. To date, the incubation period of the coronavirus (from infection to the onset of symptoms) has averaged around five days, while the average period between infection and resolution (recovery or death) has averaged around 14 days. There is some evidence of longer periods, but these analyses are clouded by the potential for a recovered individual to have been separately infected again. Combining an R-value of 2.5 with a term of illness of 14 days implies an exponential daily growth rate of infections of just over 9 percent. A daily growth rate of less than 5 percent implies an R-value of less than one, meaning that the epidemic would eventually die out.

The second quantity is the number of people who are infected and knowingly continue to interact with uninfected persons, essentially breaking quarantine. Several published medical studies on quarantine compliance indicate that the effectiveness of isolation ceases to materially improve once compliance is at 90 percent or more. One prominent study of the coronavirus progression done by Imperial College[6] assumed a United Kingdom compliance rate of only 75 percent, which unsurprisingly led to very high estimates of infection and related mortality.

On March 3, the World Health Organization published an assumption of mortality rates at 3.4 percent of known cases of infection. As noted above, this estimate is probably much too high for the entire world. The rate of mortality has had wide cross-sectional variation across countries. This is likely to be the result of different prevalence of underlying medical conditions (heart disease, diabetes) and smoking, as well as the availability of high-quality medical treatment.

As of February 23, the CDC worst-case estimate for the United States was 1.7 million deaths, about 0.5 percent of the total U.S. population. By early May, the recorded mortality rate was about 8 percent of known cases, and it was reported to be as high as 15 percent in some countries (as of May 18, 2020, the U.S. rate was around 6 percent, with a corresponding global rate of 6.5 percent). Other studies[7] assert a much lower mortality rate (less than 1 percent). The wide disparity arises from the uncertainty of how many people have been infected but have not had positive diagnoses, either from being unaware or untested. If the lower mortality rates turn out to be correct, our model implies that the portion of "infected but unaware" persons is much larger than we might otherwise conclude.

A particularly interesting example of the variation in mortality rates is the situation of the cruise ship Diamond Princess, which was held at sea for about two weeks with an unknown number of infections already existent. After disembarkation, 712 passengers and crew were diagnosed as positive for coronavirus, a rate of about 19 percent of the 3,700 persons on board. Of those 712 infected persons, 13 deaths resulted, for a mortality rate of around 2 percent. This was a somewhat surprising result given the typical age distribution aboard cruise ships is heavily weighted toward older passengers.

One factor that was often cited as having the potential to greatly increase the reported mortality rate (i.e., 8 percent to 14 percent) is that if the infections become sufficiently widespread, hospitals would run out of the capacity to care for infected patients. Certain areas of northern Italy have been widely reported as having been at hospital overcapacity for a period of time. As of early May, the percentage of infected patients in serious or critical condition is around 2.5 percent of active (known) cases globally but only 1.5 percent in the United States. Although the breaking point of hospital capacity would vary widely from country to country, the worst-case scenario is that patients who are seriously ill from coronavirus would not receive medical care and the mortality rate would rise materially. More likely, however, is that due to the contagious nature of the illness, coronavirus patients might receive preferential treatment over patients with non-communicable conditions (e.g., cancer or diabetes). As such, there could be an indirect increase in medical mortality that would not be immediately attributable to coronavirus infection.

IMPLICATIONS FOR LONG-TERM INVESTORS

The impact of the variation in mortality rates from armed conflict casualties on financial markets is presented in diBartolomeo and Hoffman (2015), which examined data from the 1890s to around 2005. Even if the implausible worst case of 60 million deaths were to come about, this would represent about a 9-percent increase in global total mortality over a 10-year investment time horizon (much smaller than multiple years of sustained large-scale war). In the armed conflict study, the correlation of decade-long global equity market returns with the variations in mortality rates was on the order of −0.35 (mortality rates up, equity market returns down).

Let's walk through the financial algebra for a typical investor. Initial assumptions are:

- Future equity returns would be 6 percent in an average year with a volatility of 15 percent (this is a geometric mean return of 4.875 percent) or a cumulative return of 59 percent over a decade
- Future fixed interest returns would be 2 percent with a volatility of 7 percent (this a geometric mean return of 1.755 percent annually)
- The correlation of equities and fixed interest returns is 0.3
- The investor is 60-percent in equities and 40-percent in fixed interest by asset value
- The total portfolio expected arithmetic return works out to 4.4 percent with a volatility of 10.67; this equates to a geometric mean return of 3.83 percent annually or a cumulative return of 45.6 percent over a decade

If an investor had a 10-year time horizon and the pandemic effects are similar to war, the expectation of the cumulative return of the portfolio would be −1.44 percent. This arises from the expected return of 45.6 percent over 10 years times the 9-percent total increase in mortality over the decade times the correlation coefficient of −0.35 from the 2015 study. The expectation of cumulative return over a decade declines from 45.6 percent to 44.16 percent, which is a net geometric mean return of 3.72 percent. The expectation of the geometric mean annual return on the investor's portfolio has declined by only a very modest one-ninth of 1 percent per annum, conditional on an extremely grim scenario for mortality.

Other considerations at play suggest this could be a pessimistic assessment in other ways as well. War destroys physical capital (factories, roads, trucks) and war is expensive to wage, heavily impacting bond markets. During World War II, the U.S. military budget reached 35 percent of gross domestic product (GDP) for a sustained period. The massive economic stimulus program just undertaken by the U.S. government represents a one-time event of about 18 percent of GDP. War casualties are also skewed toward younger persons who would otherwise be the most productive members of an industrial society. The coronavirus mortality is skewed toward the elderly. Total per capita income in the United States tends to almost equal the overall median for persons between 60 and 80, but most of this income is unearned (private pension payments, government payments, and investment income). As such, the economic impact of each death is smaller.

As physicist Niels Bohr said, "It is always hard to forecast, especially about the future." However, it seems that the crashing of financial markets around the world early in the pandemic period cannot be explained by rational actions of long-term investors conditioned on the only available data on large variations of mortality rates. We must therefore fall back onto a couple possible explanations.

Our first speculation is that investors are very short-sighted, so that nobody is thinking logically about 10-year horizons right now. Obviously, a worst-case scenario of 60 million dead over a span of a few months is very, very, scary, as are the likely economic impacts of virus mitigation efforts. Investors switching from risky assets to riskless assets (e.g., cash, sovereign debt) have a relatively simple decision because they don't have to decide what the risk-free asset is. However, investors choosing to move from riskless assets to risky assets have to make decisions about which risky assets they believe are appropriate. This means that selling is always faster than buying, which leads to crashes, which are eventually made up for by long growth periods (e.g., the Global Financial Crisis was followed up by an 11-year bull market in global equities). A similar long expansion was experienced in the Roaring 20s, the decade of prosperity that followed the end of World War I and the 1918 Flu Pandemic.

Significant research also suggests that investors are relatively indifferent to small changes in their wealth levels, but they are extremely sensitive to larger changes, so the current large amount of selling by investors may be partly logical, but mostly it's not. You can think of this in terms of the various processes defined by Wilcox (2003), Barro (2005), and Gabaix (2009). Alternatively, we can apply the behavioral framework of cumulative prospect theory, advanced by Kahneman and Tversky (1992).

Although there is always some non-zero chance that the spread of the virus could spin out of control in heavily populated countries such as India and Brazil, we have not yet seen massive numbers of infections so far. Given the apparently reduced likelihood of that kind of worst-case situation, we assert that long-term investors should be impacted minimally by the pandemic if they are globally diversified. Of course, as in war, the impact of the pandemic on individual country markets may vary widely.

INVESTOR RISKS IN THE SHORT RUN

One explanation for the large recent decline in equity markets is that investors are not thinking long term; rather, they are thinking about their portfolio assets in a "one day at a time" fashion. Although we take no position on the likely success or failure of measures taken by various governments to stimulate economic activity and calm investors, we can certainly shed light on daily variations in investor confidence.

To assess very short-term risk, we have maintained our U.S. Short Term risk model since 1997. In this approach, a statistical factor model is adjusted daily for changes in the implied volatility of options traded on equities in the United States. A mathematical process then maps the day-to-day changes in security level volatility across the factors of covariance so that changes in market conditions are applied to all securities, not just equities in which options are traded. Security coverage includes all non-U.S. equities traded on U.S. exchanges in American depositary receipt (ADR) form, so the model does cover most large publicly traded firms globally. Mathematical details of the model are presented in diBartolomeo and Warrick (2005). Although volatility levels are presented in the usual annualized units, the intended time horizon for the risk forecast is the next trading day. This estimate is different from the traded VIX contract because the volatility estimate for the S&P 500 is built from the estimated volatility and weights of the individual securities comprising the index.

Figure 1 shows the expected volatility of the S&P 500 on an equal-weighted basis, the S&P 500 on the conventional capitalization-weighted basis, and the tracking error of the two portfolios. The time period is from the end of October 2019 through to March 16, 2020, at a daily frequency.

Figure 1 ANNUALIZED VOLATILITY

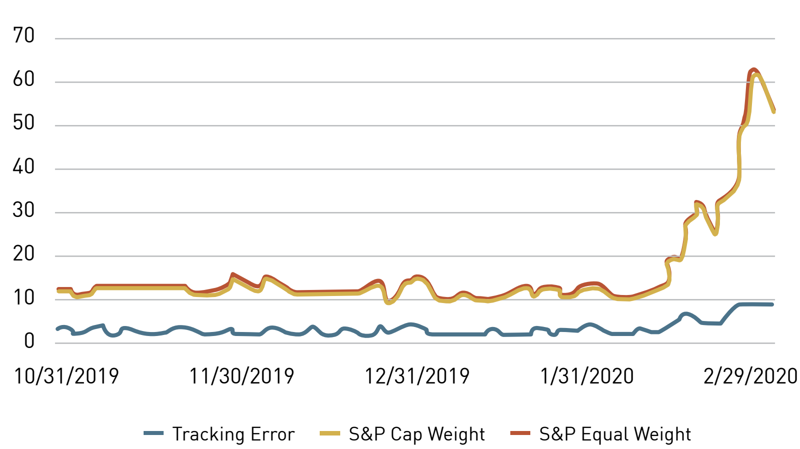

Figure 1 shows that the volatility values for both series increased roughly five-fold over the sample period. Peak annualized volatility values of more than 60 percent were observed on March 13 but decreased to 52 percent for the cap-weighted index and 55 percent on the equal-weighted index by March 16. As is common in crisis periods, correlations had increased with the expected correlation of the two portfolios going from 0.968 at the start to 0.984 at the end. As of the close of trading on April 15, this value was 39.8 percent.

INFERENCE ON THE LENGTH OF THE PANDEMIC PERIOD

We can use the above information to make a useful inference about U.S. equity investors' expectations of the persistence of heightened risk from the pandemic. Let's assume that at the start of the period (November 2019), investors believed the expected total return of the S&P 500 was 6 percent with a known dividend yield of 2 percent, implying a growth rate of 4 percent. We also will assume that the market was fairly valued, implying that expected returns and required returns were equal.

During the sample period, the expected volatility of the S&P 500 increased by roughly 40 percent, and the index level declined roughly 20 percent by mid-March. We assert that rational investors would respond by increasing their required rate of return by 6.7 percent per annum (40/6) for the period of heightened risk. The denominator scalar of six is derived from the implied risk boundary of an investor's original risk level.

We next utilize a simple Gordon dividend discount model (Gordon and Shapiro 1956) for the pricing of a stream of $1 dividends.

$$P\text{ (November 2019)} = 1 / (0.06 - 0.04) = \$50$$

This gives us a result of $50 per $1 of dividends (2-percent yield). If we now add an increment to the required return for the heightened expected risk, we get:

P (March 2020) = 1 / (0.127 – 0.04) = $11.49

So, if investors believed that the massively heightened risks would be permanent, the S&P 500 should have fallen by 78 percent, not the observed roughly 20 percent to mid-March 2020. To get the 20-percent drop to make sense, the long-term required rate of annual return must increase from 6 percent to 6.5 percent. Using our "rule of six," this means that investors are pricing the market as if the long-term expected volatility increased by 3 percent (e.g., from 12 percent to 15 percent). Updating our parameters to mid-April 2020, the expected decline of the S&P 500 should have been 58 percent for a permanent risk increase, but the observed decline was roughly 9 percent.

If we assume that the median survival time for a publicly traded firm in the United States is about 20 years as asserted in diBartolomeo (2010), we can infer how long investors expected the heightened risks to last. The underlying model in this paper is based on the contingent claims concept from Merton (1974). Essentially it argues that stockholders have two options that lenders don't have. One is a call option on the assets of the firm that can be exercised by paying off the firm's debt. The other option is a put option associated with the limited legal liability of shareholders. If the assets of a firm fall sufficiently, the shareholders can walk away, transferring the assets to lenders. The "underlying" of these options is the assets of the firm, and the strike price is approximately the value of the firm's debt. Using option pricing models, we can solve the relationships for the expected expiration date of the options, which is the median of the expected distribution of firm survival time.

Using the mathematical property that variances are additive, we obtain an implied length for the pandemic of approximately seven months as of mid-March and five months as of mid-April. Please note that this analysis ignores the increase in the S&P dividend yield over the sample period from 2 percent to about 2.4 percent, and the potential loss of a modest amount of total return (–0.11 percent per annum) over long horizons as previously discussed. This result suggests that investors remain conservatively biased, as the University of Washington IHME pandemic model does not project actual disease impact beyond August 4 in the United States.[8]

A RETURN TO NORMALCY?

The severity of the pandemic appears to have peaked in many countries, including the United States and most of Western Europe. As such, there is extensive debate within countries about how and when to reduce various forms of lockdown restrictions on their populations so as to minimize the already great damage to their respective economies. We will assume that no country would undertake to reduce restrictions unless its leaders believed the worst was over and that the number of future infections would be manageable within their healthcare systems.

The question at hand is whether a particular nation can be confident in their views. To address the issue of "assumption dependence," we will consider two metrics of readiness for a range of nations. The first is the number of coronavirus tests performed to date as a percentage of the population. The second is the percentage of tests that have given a positive result for infection. It should be noted that "tests given as a percentage of the population" does not imply that testing is that widespread in the general case (e.g., healthcare workers may have been tested multiple times).

We believe an even more important metric of readiness to return to normalcy is the percentage of positives among tests given. Given that the supply of tests is limited, medical authorities would naturally allocate tests to people most likely to need testing (those showing relevant symptoms), where a high percentage of positives is to be expected. Testing of random samples of the general population is likely to show much lower percentages of positive outcomes.

As of mid-April, the worst percentage was in France with more than 44-percent positives, followed by Brazil at 40.7 percent. It should be noted that New York State was also over 40 percent. The lowest was in the United Arab Emirates with just 0.7 percent of tests producing positives. Of the 14 nations with more than 20,000 infections, the average national rate was 20.75 percent, with the United States and United Kingdom at 19.75 percent and 26.2 percent, respectively. Many more rural U.S. states also reported lower percentages of positives (e.g., Vermont reported 7.1 percent).

Given the high coefficient of variation across even neighboring countries and regions, the process of coordinating a gradual return to normalcy may be prolonged as a matter of the "lowest common denominator." It would seem that an increase in the mass production of testing kits will be a key determinant of the date at which the world can declare the battle with coronavirus is at least a stalemate, if not a victory.

CORPORATE CREDIT RISK

An obvious area of investor concern during the current coronavirus pandemic is the ability of corporations to make timely payments on their bond debt and bank loans. We begin with a very familiar measure, the yield spread between investment-grade corporate bonds denominated in U.S. dollars and U.S. Treasury bonds of similar maturity as shown in figure 2. This incremental yield is compensation to investors for both the risk of default and the lower liquidity of corporate bonds. The sample period is from the beginning of 1992 to near the current date.

Figure 2

YIELD SPREAD BETWEEN INVESTMENT-GRADE CORPORATE BONDS

Denominated in U.S. dollars and U.S. Treasury bonds, 1992–2020

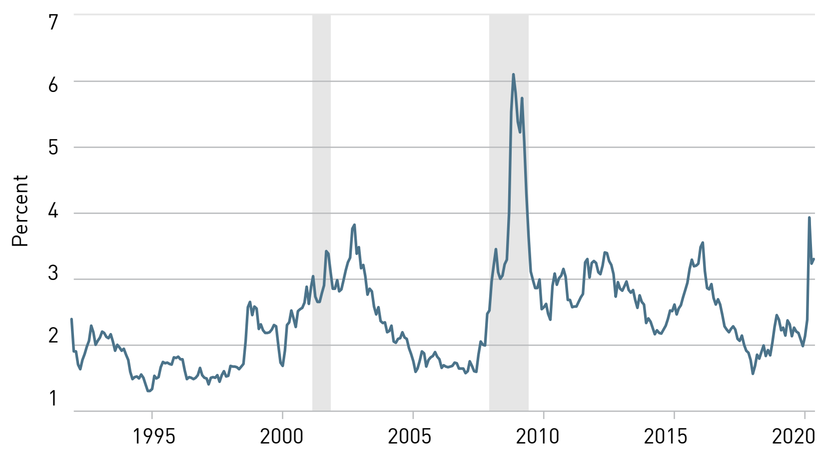

Source: Federal Reserve Economic Data (FRED), fred.stlouisfed.org

Effectively, equity becomes a portfolio of a call option and a put option on the assets of the firm with a strike price approximately equal to the value of the firm's debt. In diBartolomeo (2010), the expiration date of the options was solved numerically as the implied central tendency of the distribution of the survival time of the firm.

The cited model is quite complex, but we can illustrate the concept in a simple fashion. As long as shareholders can issue new equity to get cash to pay off debt, a firm should be able to survive and avoid default. Therefore, the ratio of the debt value to the market capitalization of a firm should be a crucial metric. If the debt is small compared to the market value of equity, it should be easy for a firm to avoid default. If the debt is large compared to the market value of equity, the firm may have difficulty raising enough cash to pay off the debt in the event of operational difficulties arising from the coronavirus pandemic.

Figure 3 provides a scatter plot of the relationship between the debt-capitalization ratio for a universe of equities and corporate bond yield spreads. To avoid any seasonality issues, each data point is defined as of March 31 of a particular year, starting in 1992. The ratio is the simple average of all stocks publicly traded in the United States at each date (including ADRs) with market capitalization of more than $250 million.

It should be noted that a significant number of the data points lie to the right of the March 31, 2020, data, indicating that even less favorable credit conditions have prevailed over material portions of the sample period. In short, although the stock market had come down a lot, this decline was from a high level, so the month-end status of this measure is better than about a quarter of previous observations. The Pearson correlation of the two series is 72 percent and the R-squared is 0.52, which is highly statistically significant, given the size of the data sample.

The most recent month-end data for March 31, 2020, is highlighted in red. This point is approximately 1.15 percent above a simple ordinary least squares (OLS) trend line. One can rationalize this extra yield with two closely related but not identical explanations. The first rationale is that there is currently a heightened degree of uncertainty in terms of what future default rates and recoveries may be, and investors demand additional yield to compensate them for this uncertainty. The functional form is presented in diBartolomeo (2018, slide 15).[9]

Alternatively, investors may believe that the coronavirus pandemic represents a circumstance where the correlation of likely defaults across many borrowers is likely to be very high for at least a short period. This means that the distribution of return outcomes on diversified portfolios will act more like a single bond and exhibit negative skew and positive excess kurtosis (i.e., a typical bond has a lot more room to go down in value than up). Using the same method again, the expectations of higher volatility would require higher returns (i.e., spreads), as described above.

We can also integrate the above ratio analysis with the updated dataset on corporate sustainability. For the 26-year sample from 1992–2017, the simple average of "expected half-life" (50-percent probability of surviving this long) is 17 years across time for our equity universe of U.S. firms (including those with less than $250 million in capitalization). This longevity value changes materially if firms are weighted by either market capitalization (higher) or by revenues (shorter). The latter effect can be interpreted as "too big to fail" firms taking on higher financial leverage.

This average longevity estimate corresponds to a market-implied annual extinction rate of 4 percent. If we run an OLS regression of the debt-capitalization ratio series on the expected life data, we obtain an estimated half-life of around 13 years, or an extinction rate of 5.2 percent as of March 31, 2020. In the first half of April, U.S. equity market averages rose about 10 percent, decreasing the ratio of the value of debt to market capitalization to a ratio of 0.321 as of April 15, very close to the long-term sample period mean of 0.325 with a corresponding implied extinction rate of 4.1 percent, just slightly above the historic average level.

We should consider several other items when assessing bond credit risks during the pandemic. The first one is the average maturity of the debt. If debt is only a fraction of the market value of equity, the interest payments are likely to be a very small fraction of market capitalization. As of March 31, 2020, the average maturity of debt across the nearly 32,000 global issuers covered in our internal database is almost exactly five years.

At the 5.2-percent extinction rate, about 76 percent of bonds would not default within five years. The long-term average spread for the sample period is 2.45 percent, so the current spread provides about a 1.5-percent annual yield increment or about 8-percent extra return cumulatively over five years. At the 4-percent long-term extinction rate, about 81 percent of firms are expected to survive at least five years. Under current conditions, investors are being compensated 8 percent over five years for the risk of having an incremental 5 percent of firms default. There must be a risk premium involved because even if we experienced 5-percent incremental defaults with zero recoveries, the incremental cumulative loss of principal could not be greater than 5 percent. If we assumed recoveries averaging 40 percent on corporate bonds (a typical assumption), the incremental loss of principal would be 3 percent of our portfolio. This additional return (8>3) is therefore consistent with our risk premium hypothesis for the very high yield spread as of the end of March.

It also should be considered that although yield spreads have increased, the overall interest rate being paid by borrowers on new debt has not changed very much. The spread values used herein are increments in yield over U.S. Treasury 10-year bond rates, which have fallen to historically low levels of around 1 percent, compared to a long-term average of 4.2 percent. Some yields actually have fallen, particularly on corporate bank loans (where recoveries average around 80 percent). The net effect may be that the long-term

Figure 3

RELATIONSHIP BETWEEN THE DEBT-CAPITALIZATION RATIO FOR A UNIVERSE OF EQUITIES AND CORPORATE BOND YIELD SPREADS

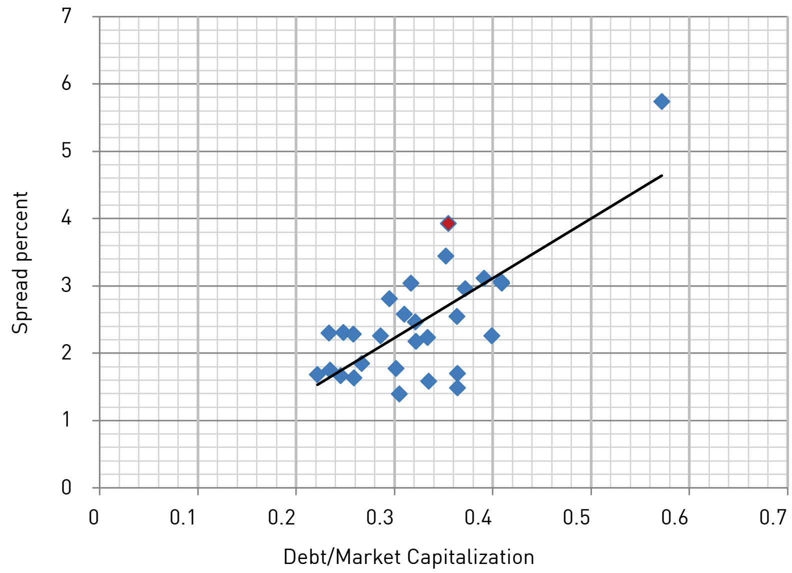

expectation of corporate sustainability would perversely improve somewhat as the annual carrying cost of debt capital declines, thereby improving corporate profits in the future.

REGULATORY IMPACT: BURYING YOUR HEAD IN THE SAND?

A recent Bloomberg article citing Johannes Borgen[10] is an interesting example of regulatory impact. The author discusses recent pronouncements by the Swiss Financial Market Supervisory Authority (FINMA). The essence of the FINMA position was that because the pandemic was not foreseeable in a traditional bank Value-at-Risk (VaR) model, any realized losses that exceeded the model expectations for losses would be ignored for the purpose of ongoing bank capital requirements. In essence, the FINMA position is that, under the regulatory specifications, the models were right but reality was wrong, so risk model failures need not be penalized in bank operations.

The perverse logic of this argument relies on the deficient mathematical properties of VaR as a risk measure. If you are measuring losses on a daily basis, the likelihood of an extreme negative event occurring on a particular day is extremely small, so it follows that at the 95-percent or 99-percent confidence interval massive potential losses can be ignored. This is like living under a live volcano and saying the risks are low because it is unlikely that the volcano will erupt today. Of course, what matters is not the daily likelihood of an extinction event but rather the cumulative likelihood of the event over as many days as you are exposed to the source of danger.

As a concept, risk is always in the future. Volatility of asset values has occurred in the past, and volatility of asset values will occur in the future, but there is no risk in the past. What is done is done. The problem is that most organizations equate past volatility and future risk in wholly unsound ways. The potential for rare extreme events of any kind is routinely ignored by financial institutions. This critical omission and how to address it has been pointed out in articles such as diBartolomeo (2018).

A simple example should suffice to illustrate the point. Let's consider the period from 1900 to the present date, a period of 120 years. During that period, at least five global events could have entirely destabilized financial markets (World War I, 1918 Flu Pandemic, 1929 Crash and Depression, World War II, Global Financial Crisis), or a probability of some kind of large event of around 4 percent per annum.

Let's assume that an investor holds a portfolio with volatility at 10 percent under typical conditions, but due to the potential for rare, large events there is a 2-percent annual probability of a tail event with a 90-percent loss. Using the method of Cornish and Fisher (1938), we obtain the economically equivalent annual portfolio volatility, which increases from 10 percent to more than 33 percent. Using such a process, even a parametric VaR estimate can reflect reality appropriately.

Regulators may choose to take the view that banking risk is really about what happens in the ordinary course of business, and that extreme events will be addressed separately through the intervention of government. We would argue that it is ill-advised for investors to assume a similar posture, and therefore we must make the effort to incorporate the potential for rare but extreme events into risk assessments.

CONCLUSIONS

It should be clear that the coronavirus pandemic is a profoundly serious matter. Allowed to spread entirely without limit, the death toll would be massive and comparable to the peak year of World War I or World War II as a percentage of the population. However, it is also obvious that steps have been taken already to slow the spread of the infection, even if those efforts have seemed to be lagging in some countries. We have presented a simple algebraic model with which an investor can track the progress of the pandemic.

Although we do not take a position on the matters of morality, ethics, or public policy associated with the pandemic, it is clear that even in the worst-case scenarios for related mortality and healthcare expenditures, the economic impact for financial market investors should be minimal when observed over long horizons of 10 years or more. For investors, the coronavirus pandemic itself is not the end of the world. Our current maximum likelihood estimate for mortality is in the range of 300,000 to 400,000, comparable to the death toll of the 2005 Asian tsunami.

Much greater risk to investors and the world economy at large comes from clumsy handling of mitigation measures than from the virus itself. Of greatest concern is that many organizations, companies, and governments have been engaged in an undesirable game of managing the optics of the situation. Whether out of concern for liability in litigation, the ability to enforce business interruption insurance, political expediency, or reputational issues, the actions of many entities are being driven by a desire to be perceived as managing risk as opposed to actually managing risk in response to factual information and analysis.

Dan diBartolomeo is president and founder of Northfield Information Services, Inc., which develops quantitative models of financial markets. He is a former faculty member at Brunel University and is the co-editor of the Journal of Asset Management. *Contact him at dan@northinfo.com.*

ENDNOTES

1. Unless otherwise noted, all statistical data on the coronavirus pandemic is drawn from the Worldometers website, https://www.worldometers.info/coronavirus/. The author is indebted to Worldometers for this extraordinary public resource.
2. Sherry Fink, "Worst-Case Estimates for U.S. Coronavirus Deaths," *New York Times* (March 13, 2020), https://www.nytimes.com/2020/03/13/us/coronavirus-deaths-estimate.html.
3. Refer to http://covid19.healthdata.org/united-states-of-america.
4. For details, see appendix 1 in Dan DiBartolomeo, "The Corona Virus Emergency: A Quantitative View," https://www.northinfo.com/Documents/937.pdf.
5. Ibid.
6. Regarding the Imperial College coronavirus study, see Mark Landler and Stephen Castle, "Behind the Virus Report That Jarred the U.S. and the U.K. to Action," *New York Times* (March 17, 2020, updated April 2, 2020), https://www.nytimes.com/2020/03/17/world/europe/coronavirus-imperial-college-johnson.html.
7. See "Coronavirus Cases Surge in U.S. and Europe," *New York Times* (March 5, 2020, updated March 17, 2020), https://www.nytimes.com/2020/03/05/world/coronavirus-live-updates.html, regarding the work of Dr. Andrew Kucharski.
8. Since this article was written, the IHME models have been updated to project the impact of the elimination of lockdowns and other social distancing measures. The IHME models still do not make any explicit projection beyond August 4, but they do project that a declining rate of related deaths will still have an expected value of around 150 per day in the United States at that date.
9. See slide 15 in Dan DiBartolomeo, "Dissecting 'Duration Times Spread,'" Northfield (2018), https://www.northinfo.com/documents/923.pdf.
10. See Matt Levine, "Money Stuff: Banks Pretend 2020 Never Happened," *Bloomberg* (April 16, 2020), https://www.bloomberg.com/news/newsletters/2020-04-16/money-stuff-banks-pretend-2020-never-happened.

REFERENCES

Barro, Robert. 2005. Rare Events and the Equity Premium. NBER Working Paper No. 11310 (May).

Cornish, E. A., and R. A. Fisher. 1938. Moments and Cumulants in the Specification of Distributions. *Review of the International Statistical Institute* 5, no. 4 (January): 307–320.

diBartolomeo, Dan. 2010. Equity Risk, Credit Risk, Default Correlation, and Corporate Sustainability. *Journal of Investing* 19, no. 4 (winter): 128–133.

———. 2018. Why Getting Risk Right is Wrong! *Northfield Newsletter* (April). https://www.northinfo.com/Documents/802.pdf.

diBartolomeo, Dan, and Howard Hoffman. 2015. Seeing the Big Picture: How Conflict and Corruption Impact Financial Markets. Northfield Working Paper (June). https://www.northinfo.com/Documents/646.pdf.

diBartolomeo, Dan, and Sandy Warrick. 2005. "Making Covariance-based Portfolio Risk Models Sensitive to the Rate at which Markets Reflect New Information," in J. Knight and S. Satchell (eds.), *Linear Factor Models in Finance*. Oxford, UK: Butterworth-Heinemann, p. 249–259.

Gabaix, Xavier. 2012. Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance. *Quarterly Journal of Economics* 127, no. 2 (May): 645–700.

Gordon, Myron J., and Eli Shapiro. 1956. Capital Equipment Analysis: The Required Rate of Profit. *Management Science* 3, no. 1 (October): 102–110.

Kahneman, Daniel, and Amos Tversky. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty* 5, no. 4 (October): 297–323.

Merton, Robert. 1974. On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance* 29, no. 2 (May): 449–470.

Wilcox, Jarrod. 2003. Harry Markowitz and the Discretionary Wealth Hypothesis. *Journal of Portfolio Management* 29, no. 3 (spring): 58–65.

5619 DTC Parkway, Suite 500
Greenwood Village, CO 80111
Phone: +1 303-770-3377
Fax: +1 303-770-1812
www.investmentsandwealth.org