Re-Examining the Credit Premium

Jared Kizer, CFA

March 22, 2018

Abstract

Using the Ibbotson/Sinquefield data documenting the returns of long-term government and corporate bonds, Asvanunt and Richardson [2017] find a sizable investment-grade credit premium that is also statistically significant after accounting for exposure to equity, size, value and momentum factors. In this piece, I focus on the Ibbotson/Sinquefield data starting in 1969 when the series began directly referencing the Salomon Brothers High-Grade Long-Term Corporate Bond Index spliced with Bloomberg/Barclay's data starting in August 1988 and separately the Bloomberg/Barclay's data going back to 1973. For both series, I find a weaker credit premium that is *not* statistically significant after adjusting for equity market factors.

The differences in findings appear to be driven by the abnormally high investment-grade credit (IGC) premium associated with the early history of the Ibbotson/Sinquefield data. An examination of that early history calls into question whether the data is truly representative of investment-grade corporate bond returns net of defaults and downgrades. The primary results questioning the early period data are 1) a significantly higher Sharpe Ratio for the IGC premium compared to either equities or interest rate risk during a period encompassing the Great Depression 2) anomalously high returns during the 30s, which was a period of significant credit stress 3) an IGC premium in excess of 2.5 percent per year in this early period for an index focused on Aa and Aaa corporate bonds and 4) most significantly, an IGC premium embedded in the Ibbotson/Sinquefield data that seems completely disconnected from the frequency of corporate bond defaults reported by Moody's during common periods of time. Consequently, I argue that the two samples I focus on are likely more representative of the true returns associated with the IGC premium while still retaining the advantages of a longer-term sample, which motivated the work by Asvanunt and Richardson.

Please send any comments to jkizer@bamadvisor.com. Thanks to Corey Hendershot for research support. This analysis is for academic purposes only. The research, opinions and data shared within this paper are those of Mr. Kizer and do not directly reflect those of Buckingham Asset Management, LLC.

1 Introduction

Asvanunt and Richardson [2017] document the long-run returns and risk associated with the investment-grade credit premium (henceforth "IGC premium"). Their two main findings are 1) that the IGC premium has been sizable and statistically significant over the long term and 2) that the IGC premium has been positive and statistically significant even after adjusting for equity market beta and other equity factor exposures. The fact that the IGC premium has been positive is not surprising but it's not clear whether one should have expected the IGC premium to be positive and statistically significant after adjusting for equity market factors given the theoretical relationship between equity market risk and credit risk (see Merton [1974]).

In this paper, I refocus on both findings using two data sets. For the first series, similar to Asvanunt and Richardson, I splice the IGC premium estimated using the Ibbotson/Sinquefield (henceforth Ibbotson) data beginning in January 1969 with the IGC premium reported by Bloomberg/Barclay's (henceforth Bloomberg) starting in August 1988. The second series uses only Bloomberg data and estimates the IGC premium using this data for the period of December 1974 through July 1988 and splices this series with the same IGC premium used in the first series that Bloomberg began calculating in August 1988. For both series I find a weaker credit premium that is *not* statistically significant after adjusting for equity market factors.

Relative to Asvanunt and Richardson, the differences in my findings appear to be driven by the abnormally high IGC premium associated with the early history of the Ibbotson data. An examination of that early history calls into question whether the data is truly representative of investment-grade corporate bond returns actually earned during that historical period. The primary results questioning the early period data are 1) a significantly higher Sharpe Ratio for the IGC premium compared to either equities or interest rate risk during a period encompassing the Great Depression 2) anomalously high returns during the 30s, which was a period of significant credit stress 3) an IGC premium in excess of 2.5 percent per year in this early period for an index focused on Aa and Aaa corporate bonds and 4) most significantly, an IGC premium embedded in the Ibbotson data that seems completely disconnected from the frequency of corporate bond defaults reported by Moody's during common periods of time.

Section II outlines the data sources used in this paper. Section III examines the characteristics of the IGC premium estimated using the pre-1969 Ibbotson data relative to other premia over the same period of time. Section IV analyzes the IGC premium using the two series I constructed and described above. Section V analyzes whether portfolios that exclude corporate bonds can reasonably replicate the returns of corporate bonds. Section VI discusses the practical implications of my findings and Section VII concludes.

2 Data Sources

I use Ibbotson returns on long-term government bonds and long-term corporate bonds covering the period of January 1926 through December 1969. These returns are obtained from DFA's Returns 2.0. program. From Bloomberg I use the excess returns associated with the U.S. Corporate Index from August 1988 through October 2017. I also use the total returns reported by Bloomberg for the U.S. Corporate, U.S. Treasury and U.S. Long-Term Treasury indexes for the period of January 1973 through October 2017. From AQR's data library, I use their "CORP XS" return data for the period of January 1969 through July 1988. This series is created from the Ibbotson long-term government and corporate bond data as described in Asvanunt and Richardson [2017]. From Ken French's data library I use the monthly returns associated with the U.S. market, size, value and momentum premia as well as U.S. small-cap value equities and the risk-free rate for the period covering January 1969 through October 2017 and annual returns for the U.S. market premium and risk-free rate covering 1930 through 2009. I reference Moody's default data from their 2016 study on corporate defaults and recovery rates. This study is included in the references section.

3 Examining the Pre-1969 Ibbotson Data

3.1 Risk and Return

The pricing data used by corporate bond market index series is generally known to have quality issues, namely that pricing might not be indicative of tradable levels due to the illiquidity and heterogeneous nature of the market. These issues alone might make one question how reliable data from the 30s, 40s and 50s could possibly be but here I focus on more quantifiable concerns working directly with the Ibbotson data from these earlier periods.

I specifically focus the analysis on the pre-1969 period because as Ibbotson and Sinquefield [1976] note 1969 was the year in which the Ibbotson data began directly referencing the Salomon Brothers High Grade Long-Term Corporate Bond Index whereas during a significant fraction of the earlier period Ibbotson and Sinquefield attempted to replicate what the returns of the Salomon Brothers index might have been had it existed.[1] Given this issue and the general concerns about early/mid 20th century data quality, examining the pre-1969 returns data seems like a worthwhile exercise.

The first step in the analysis uses 24-month rolling regression analysis and the Ibbotson long-term government and long-term corporate bond returns to create a government bond return series that is approximately duration neutral with respect to the long-term corporate bond return series. I use "approximately" because there is no way to insure exact duration neutrality within a regression analysis framework but this method is certainly a reasonable and common approach when duration data are not explicitly available. The rolling regression equation is as follows:

$$LTC_t \;=\; \alpha \;+\; \beta \times LTG_t \;+\; \epsilon_t$$

The appropriate βs from the rolling regressions are then multiplied by the Ibbotson long-term government bond return for a particular month to create the new government bond return series. These monthly returns are then used to create annual returns. The second step is taking the difference of the annual returns of the Ibbotson long-term corporate series and the new government bond series (labeled LTG_{Neut} below) to create the IGC premium series.[2]

$$IGC\ Premium \;=\; LTC - LTG_{Neut}$$

As in Asvanunt and Richardson [2017] the idea here is to create a measure of the IGC premium that is not contaminated by either a long or short position that co-moves with general interest

[1] Note, however, that the Ibbotson series referenced the S&P High-Grade Corporate Composite Bond Index during the 1926–1945 period.

[2] The reason I chose not to work directly with the IGC premium posted in AQR's data library for some portions of my analysis is the need to build an annual IGC premium return series, which cannot be done accurately by simply compounding AQR's monthly IGC premium return series.

rate changes. My methodology produces annual returns for the IGC premium over the period of 1928–2016 but the analysis in this section focuses on the 1930–1969 period since the data used by Asvanunt and Richardson starts in the mid 1930s. Table 1 reports the annualized returns, volatilities and Sharpe Ratios of the market, term and IGC premia over the 1930–1968 period.

Table 1 shows that the Sharpe Ratio for the IGC premium was over 2x higher than the market premium and almost 5x higher than the term premium over the pre-1969 period. It is also worth noting that the IGC premium itself was 2.6 percent per year. While this might not seem high one has to keep in mind that the Ibbotson series is apparently oriented toward the highest grade corporate credits over its early history (see Hallerbach and Houweling [2013]). An IGC premium of 2.6 percent per year for investing in Aa and Aaa corporate bonds seems high. I now turn to results from specific sub-periods of the early period data. Figure 1 reports the decade-by-decade Sharpe Ratios for the market, term and IGC premia for the 30s, 40s, 50s and 60s.

Notably here we see two decades, the 30s and 40s, where the Sharpe Ratio of the IGC premium was well in excess of 1.00. In examining the decade-by-decade history of the market and term premia from the 30s through the first decade of the 2000s there was not a single decade where the Sharpe Ratio of either was 1.00 or higher. Further, of the four early period decades the IGC premium had the highest risk-adjusted returns in three.

The result from the 30s is particularly suspicious. This was a decade with massively negative returns for equities in a handful of years and consequently a relatively low realized Sharpe Ratio for the market premium. Nevertheless, the Sharpe Ratio for the IGC premium was 1.66, 12x higher than the Sharpe Ratio for the market premium and 2x higher than the Sharpe Ratio for the term premium. This strains credibility, particularly when also noting that the compound return of the IGC premium was 5.5 percent per year while the compound returns of the market and term premia were –0.9 percent and 4.3 percent per year, respectively. It's hard to imagine how either the Sharpe Ratio or the compound return of the IGC premium could have truly been this high relative to the other two premia in a decade highlighted by the Great Depression. The result from the 40s is only slightly less suspicious given it resulted in a Sharpe Ratio in excess of 1.40. Next I examine the year-by-year returns of the IGC premium. Figure 2 shows this data.

While the compound return and Sharpe Ratio summarize periods and sub-periods of the data, visually inspecting the year-by-year returns data shows that the IGC premium was positive in virtually every single year of the pre-1969 period. In fact, it was positive for *17 straight years* from 1930 through 1946. This level of consistency again makes one question whether this data is truly reflective of the returns associated with credit risk. Looking at individual year returns the IGC premium was +7.8 percent in 1930 when the market premium was –31.2 percent and +2.8 percent in 1937 when the market premium was –35.0 percent as two examples of many that are hard to reconcile. These results are also depressing the market beta of the IGC premium below what it otherwise would be, leading to an understatement of the market risk embedded in the IGC premium and an overstatement of the returns of the IGC premium net of equity market factors.

3.2 Moody's Corporate Bond Default Data

Moody's has data that provides a sanity check of the Ibbotson data. Each year Moody's releases its "Annual Default Study" of corporate bond defaults going back to 1920. While I have documented concerns with the early period returns data, it's likely that data simply tracking whether an issuer defaulted or not is more reliable. This data is also likely indirectly indicative of periods of relatively higher and lower amounts of downgrades since one would expect a period of higher numbers of defaults to be positively correlated with a period of higher numbers of downgrades. Both defaults and downgrades are events that would tend to depress the returns of the IGC premium in a particular period, at least to the extent that these events were surprises to the market. Figures 3 and 4 graph the issuer default counts Moody's reports on investment grade and high-yield bonds, respectively, over the period of 1920–2016.

While what Moody's defines to be an investment-grade rated security or a speculative/high-yield security may have changed over time (for example, most financial market histories point to the 1980s as the origination of the high-yield market as we know it today), the Moody's data shows that there were a significant number of investment-grade and speculative-grade defaults in the early part of the sample. This seems to conflict with the IGC premium observed in the 30s even ignoring the additional impact of downgrades. It's also interesting to note, however, that there were virtually no investment grade defaults reported by Moody's over a 20+ year horizon from the early 40s through the mid 60s. This result is also odd, however, in that the highest returns for the IGC premium were in the 30s compared to any of the other earlier period decades. Would one not have expected higher returns in the 40s, for example, where realized defaults appear to have been low following a period of extraordinarily high defaults?

Digging further into this relationship Figures 5 and 6 report scatter plots of investment grade issuer default count vs. the annual IGC premium in Figure 5 and the investment grade issuer weighted default rate vs. the annual IGC premium in Figure 6. One would expect the slope of this relationship to be negative i.e., years with higher default counts or higher default rates would correpond with a lower or negative IGC premium, but this is not what we see. We in fact see a *positive correlation* between both measures of default and the IGC premium, indicating bizarrely that the IGC premium tended to be higher in years with higher default counts or default rates. While all the above results damage the credibility of the early history of the Ibbotson series it could be argued this result is the most damaging.

3.3 What Might be Driving Data Quality Issues?

IGC premium data quality issues can be placed into one of three categories. First, we know that corporate bond data is suspect due to the de-centralized nature of trading and the relative illiquidity of the market. For example, while equities have historically traded on centralized exchanges where closing prices are observable, this is not true for corporate bonds. Further, while each company generally has one class of equity a given company can and most commonly does have multiple bonds trading in the marketplace with different liquidity, maturities and so on. The equity of a given company is also generally more liquid than any bonds it has issued. Related to the Ibbotson series the argument here would be that data quality on corporate bond pricing was so bad in the early history of the Ibbotson data that the end result, the Ibbotson long-term corporate bond series, is simply not representative of actual corporate bond market pricing during most or all of this early period.

Second, from Ibbotson and Sinquefield [1976] we know a portion of the early period data was an attempted replication of what the returns of the Salomon Brothers High-Grade Long-Term Corporate Bond Index might have been had it actually existed. While Ibbotson and Sinquefield no doubt tried to do the best they could with what they had to work with, this attempted replication could be related to the results I've found. It could also be that the S&P High-Grade Corporate Composite Bond Index used in the earliest portions of the Ibbotson data is somehow not accurately representing corporate bond market returns for some or all of its contribution to the sample.

Third, it's not clear whether the early period data accurately accounts for downgrade and default events within high-grade corporate bonds. While defaults are rare they do occur and downgrades occur more frequently. If the total returns of the underlying Ibbotson series do not accurately account for the generally negative return events associated with either defaults or downgrades, the IGC premium as calculated would be overstated. As one example if the underlying yield/price data that Ibbotson used simply measures the yields/prices of bonds at each month end *that are still Aa or Aaa*, it would be missing the likely negative return events associated with bonds that were Aa or Aaa at the beginning of a month but were no longer Aa or Aaa (or had defaulted) at the end of that month. While there is no way for me to confirm this with any research that I have examined, I suspect this is a significant issue with at least some portions of the early period Ibbotson long-term corporate bond return data. This set of concerns might also fit with the earlier finding that the IGC premium was strangely positively correlated with annual default count.

In light of all the above, I think it's fair to question whether the early period IGC premium actually measures what it claims to measure. The Sharpe Ratio, raw returns and sub-period results all suggest the IGC premium as measured is significantly inflated over the period of 1930–1968. Further, the Moody's default data from the 30s does not align with an IGC premium in excess of 5 percent and a Sharpe Ratio in excess of 1.50 and the results from the 40s and 50s seem bizarre compared to the 30s. The right step from here would appear to be focusing on more recent period results while still retaining a long-horizon sample. To do this, I focus on Ibbotson data starting in 1969 and Bloomberg data starting in 1973.

4 Examining Later Period Data

4.1 Ibbotson Data Spliced with Bloomberg Data

I start the later period analysis with the Ibbotson IGC premium series from AQR's data library that I then splice with the excess return of the Bloomberg U.S. Corporate Index as reported by Bloomberg beginning in August 1988.[3] Using monthly returns data, Table 2 reports the annualized return, volatility, Sharpe Ratio and t-stat of the IGC premium over the period of 1/1969–10/2017.

The IGC premium achieved a compound return premium of 1.1 percent per year, a Sharpe Ratio of 0.27 and a t-stat of 1.9. Given an expectation that the premium should be positive, I think it's fair to categorize this result as statistically significant even though the t-stat is a bit shy of 2.0. Let's now look at the premium after adjusting for equity market factors.

For this spliced IGC premium series, I run nine different regressions using monthly returns data. The sets of independent variables are:

1. Market premium with no lags
2. Market premium with one lag

[3] There is no intention to claim this series is the more representative or accurate of the two series I analyze.

3. Market premium with two lags
4. Fama-French factor model with no lags
5. Fama-French factor model with one lag
6. Fama-French factor model with two lags
7. Carhart factor model with no lags
8. Carhart factor model with one lag
9. Carhart factor model with two lags

The Fama-French model adjusts for size and value exposure in addition to equity market exposure while the Carhart model adjusts for size, value and momentum exposure in addition to equity market exposure. The reason for running these regression analyses is to determine whether there's anything unique about the IGC premium that isn't accounted for by standard equity market factors (i.e., whether the regression alpha is positive and statistically significant). If so, this strengthens the argument for owning corporate bonds even in the presence of factor-tilted equities. If not, it indicates that corporate bonds may be redundant in the context of portfolios that own factor-tilted equities and government bonds. As a reminder, Asvanunt and Richardson [2017] find the former, arguing that credit risk is not entirely redundant in the presence of equity market factors. Table 3 reports the regression results for the Ibbotson IGC premium series spliced with the Bloomberg series.

Reviewing the monthly alpha estimates for each of the nine specifications, we see that none are statistically significant with the first seven generally not close to statistical significance. Note that this means the alpha estimates aren't even statistically significant in the regressions that adjust for nothing more than broad U.S. equity market exposure. While specifications eight and nine approach statistical significance this is only after accounting for the negative momentum loading apparently embedded in the IGC premium. This is an important note because the negative momentum loading would have generally detracted from historical return (hence the smaller alphas in the specifications that didn't include UMD) and this detracted return was of roughly similar monthly magnitude to the alpha that is left over. In other words to capture the alpha one has to either accept the generally negative baggage of anti-momentum exposure or find ways to own corporate bonds while avoiding the negative momentum exposure apparently associated with a generic approach to the space.

The factor loadings provide other interesting insights. Looking jointly at the MKT, SMB, HML and UMD loadings, the IGC premium looks somewhat like owning a small fraction of one's portfolio in U.S. small-cap value stocks, with specifications that include SMB and HML showing positive loadings on each. Looking at the specifications that use lagged variables we generally see those are statistically insignificant other than for UMD. So, on balance it's not clear whether the IGC premium series as constructed here suffers from stale pricing as one would expect stale pricing to result in positive and statistically significant coefficients for more of the lagged variables, particularly MKT.

It's also important to note that the R-squared ratios are generally low indicating a substantial amount of noise in the returns of the IGC premium that is not accounted for by equity market factors.

4.2 Bloomberg Data

While Bloomberg did not explicitly begin reporting the IGC premium until August 1988 they report monthly total returns on a wide variety of indexes as early as January 1973. I use the pre-August 1988 total returns data to estimate a monthly IGC premium and I then splice this data with the excess return of the Bloomberg U.S. Corporate Index as reported by Bloomberg beginning in August 1988. One advantage of this series is that it uses data from a single pricing source, Bloomberg, instead of splicing together series from multiple pricing sources, which likely changes what is being measured as the underlying pricing source changes.

To estimate the earlier period IGC premium I use a slightly different variation of the rolling regression method used in Section III. For the Bloomberg data I have, Bloomberg reports maturity but not duration for the early years of each series. Comparing the maturity reported for the Bloomberg U.S. Corporate Index to that of the Bloomberg U.S. Treasury Index one finds that for the early periods of the sample that the maturity of the corporate index greatly exceeded that of the Treasury index. Therefore, in addition to including the returns of the Bloomberg U.S. Treasury Index as a right-hand side variable I also include the returns of the Bloomberg U.S. Long-Term Treasury Index as a right-hand side variable to improve the regression fit. The following rolling regression equation is used:

$$CORP_t \;=\; \alpha \;+\; \beta_1 \times TSY_t \;+\; \beta_2 \times LTSY_t \;+\; \epsilon_t$$

I then use the βs from the rolling regressions and the monthly returns of the two Treasury indexes to create a Treasury returns series that is approximately duration neutral to the Bloomberg U.S. Corporate Index. This approach produces monthly returns for this augmented Treasury series for the period of 12/1974–7/1988. Taking the difference of the monthly returns of the Bloomberg U.S. Corporate Index and the regression-created Treasury index then produces the monthly returns of the IGC premium:

$$IGC\ Premium \;=\; CORP - TSY_{Neut}$$

Table 4 reports the annualized returns, volatility, Sharpe Ratio and t-stat of the IGC premium over the period of 12/1974–10/2017.

Here we also see a positive IGC premium but one that is notably weaker than the prior series I examined and consequently nowhere close to statistically significant. For reasonable implementation costs and manager expenses a safe assumption would be that this index has exhibited no historical credit premium. Table 5 reports the results of the same nine regression analyses done on the Ibbotson data. Not surprisingly, alpha is not statistically significant in any of the nine regression analyses. These results strongly suggest there is no credit premium that exists independent of equity market factors.

Looking at the factor loadings across the Fama-French and Carhart regressions shows results similar to those found with the prior series with the exception that the results show more evidence of stale pricing with lagged MKT coefficients generally statistically significant and positive.

4.3 Relationship to Moody's Default Data

The earlier data contained the bizarre result that the IGC premium was positively correlated with both default counts and default rates. I now examine the relationship of each of the two series covered in sections III and IV with the Moody's default data covering the same period. Figures 7 and 8 show scatter plots of the investment grade issuer default count and default rate, respectively, with the annual Ibbotson/Bloomberg IGC premium over the period 1969–2016.

Both figures show the expected negative correlation relationship between the IGC premium and default counts and rates. Figures 9 and 10 show scatter plots of the investment grade issuer default count and default rate, respectively, with the annual Bloomberg IGC premium over the period 1975–2016. These figures also show the expected negative correlation relationship between the IGC premium and default counts and rates.

The results of these four figures consequently show that the two series I constructed seem to be fair representations of what one would expect the time series of the IGC premium to be when compared to a time series of default counts and rates. In other words, these two series seem to pass the sanity test that the earlier period Ibbotson series failed.

5 Replication

Similar to the analysis done on REITs in Grover and Kizer [2017] a complementary exercise to the regression analysis above is examining whether a portfolio of assets that excludes corporate bonds does a reasonably good job of replicating the returns of corporate bonds. There will of course not be a portfolio that excludes corporate bonds that can perfectly replicate the returns of corporate bonds. But, if the replicating portfolio has similar or better risk/return characteristics than corporate bonds the basic argument is that corporate bonds are redundant. For the Ibbotson long-term corporate bond series I run two different replication analyses. Here are the two replicating portfolio choice sets for the Ibbotson series:

1. U.S. market-cap weighted equities and the Ibbotson long-term government bond series
2. U.S. market-cap weighted equities, U.S. small-cap value equities and the Ibbotson long-term government bond series

Here are the two replicating portfolio choices for the Bloomberg U.S. Corporate Index series:

1. U.S. market-cap weighted equities, Bloomberg Treasury Bond Index and Bloomberg Long-Term Treasury Bond Index
2. U.S. market-cap weighted equities, U.S. small-cap value equities, Bloomberg Treasury Bond Index and Bloomberg Long-Term Treasury Bond Index

As usual, I require portfolios weights be $\geqslant 0$ and sum to one and weights are calculated via constrained least squares regression analysis. Tables 6 and 7 summarize the statistical properties of the replicating portfolios relative to the corporate bond series. Table 6 compares the return characteristics of the Ibbotson long-term corporate bond returns series compared to the two replicating portfolios noted above while Table 7 compares the return characteristics of the returns of the Bloomberg U.S. Corporate Index to the two replicating portfolios. Note here that these are comparisons of total returns and not excess returns.

The general conclusion from each table is that corporate bonds are redundant in portfolios that own stocks and government bonds. This syncs with the regression findings in section IV. In particular, we see that the monthly average returns of the replicating portfolios are either equal to or higher than the returns of the corporate bond series. Further, the Sharpe Ratios of all four replicating portfolios are higher than the Sharpe Ratios of the respective corporate bond return series. In tandem these results indicate that diversifiable noise is all that appears to be left after accounting for the equity and interest rate risks embedded in the returns of corporate bonds. Further, it is likely that the expense profile is lower for the replicating portfolios compared to the corporate bond series.

6 Practical Implications of Findings

There are numerous practical applications associated with the results. The most obvious is that the empirical case for a strategic allocation to corporate bonds is weak. I find that the IG premium associated with both the Ibbotson/Bloomberg and Bloomberg series is well explained by equity market factors and that the total returns of both series are dominated by portfolios of stocks and government bonds.

Second, investors who do implement corporate bonds should consider the results of Ng and Phelps [2011], Israel, Palhares, and Richardson [2017] and others, which offer ideas to potentially improve the risk-adjusted returns of corporate bond portfolios since my results suggest that a generic approach to the market may yield minimal overall portfolio benefit. Third, while market timing is notoriously difficult, if corporate bonds must be owned it could be worth exploring strategies that will increase the weight in corporate bonds closer to the target allocation when spreads are relatively wide by some well calibrated measure and reduce exposure when spreads are relatively tight. As has been illustrated by equity markets, however, where valuations may now be permanently higher than they were in more distant historical periods, this strategy is susceptible to permanent/semi-permanent changes in spreads, which could render the historical data used to calibrate the strategy useless. Fourth, again if corporate bonds must be owned, relatively low manager expenses and trading cost efficiency are likely of utmost importance to insure the investor is capturing as much of the excess return as possible given there may not be much there to begin with. There are no doubt other practical implications to consider but the above are four that stand out.

7 Conclusion

The early period returns history of the Ibbotson long-term corporate bond series exhibits behavior that is hard to reconcile with the performance of other asset classes, with macroeconomic reality and with Moody's reported default statistics. Consequently, I argue that the early history of the Ibbotson series likely overstates the benefits of owning corporates bonds and that periods starting in 1969 or later are almost certainly more representative of the true risk and return characteristics associated with the long-run IG premium. Examining both Ibbotson and Bloomberg series over this more recent period, I find an IG premium that is positive but considerably weaker than the series examined in Asvanunt and Richardson [2017]. Further, neither series exhibits statistically significant alpha after adjusting for equity market factors. I also find that the risk and return characteristics of portfolios of stocks and government bonds dominate corporate bonds. My results suggest that the empirical case for owning corporates bonds is weak. For investors who must own corporate bonds, manager fees and trading expenses should be monitored closely and investors may want to consider approaches to the market that might mitigate the generally weak excess returns associated with a generic approach to corporate bonds.

References

Annual Default Study: Corporate Default and Recovery Rates, 1920-2016. Moody's Investors Service. 2017.

Asvanunt, A. and S. Richardson. "The Credit Risk Premium". *The Journal of Fixed Income*, vol. 26, no. 3 (2017), pp. 6–24.

Grover, S. and J. Kizer. "Are REITs a Distinct Asset Class?" 2017.

Hallerbach, W. G. and P. Houweling. "Ibbotsons Default Premium:Risky Data". *The Journal of Investing*, vol. 22, no. 2 (2013), pp. 95–105.

Ibbotson, R. G. and R. Sinquefield. "Stocks, Bonds, Bills, and Inflation: Year-by-Year Historical Returns". *The Journal of Business*, vol. 49, no. 1 (1976), pp. 11–47.

Israel, R., D. Palhares, and S. Richardson. "Common Factors in Corporate Bond Returns". *Forthcoming in the Journal of Investment Management* (2017).

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

Ng, K.-Y. and B. Phelps. "Capturing Credit Spread Premium". *Financial Analysts Journal*, vol. 67, no. 3 (2011), pp. 63–75.

Table 1. Risk and Return Statistics (1930–1968)

This table reports the compound returns, volatility and Sharpe Ratio associated with the annual returns of the equity, term and investment-grade credit (IGC) premiums where the IGC premium is calculated using the Ibbotson long-term corporate bond return series. See body of paper for IGC premium computational details.

	Equity	Term	IGC
Returns	10.2	1.1	2.6
Volatility	22.8	5.9	2.8
Sharpe Ratio	0.45	0.19	0.92

Table 2. Risk and Return Statistics (1/1969–10/2017)

This table reports the annualized returns, annualized volatility, annualized Sharpe Ratio and t-stat associated with the IGC premium calculated from the Ibbotson long-term corporate bond return series spliced with the excess return reported by Bloomberg starting in August 1988 for the Bloomberg U.S. Corporate Index. See body of paper for IGC premium computational details.

	IGC
Return	1.1
Volatility	4.2
Sharpe Ratio	0.27
t-stat	1.9

Table 3. Regression results (1/1969–10/2017)

This table reports the regression results associated with nine different regressions of the monthly returns of the IGC premium calculated from the Ibbotson long-term corporate bond return series spliced with the excess return reported by Bloomberg starting in August 1988 for the Bloomberg U.S. Corporate Index on various factor models with and without lagged variables. See body of paper for IGC premium computational details.

	I	II	III	IV	V	VI	VII	VIII	IX
Alpha	0.0451	0.0387	0.0414	0.0148	0.0076	0.0096	0.0635	0.0825	0.0876
	(1.0)	(0.8)	(0.9)	(0.3)	(0.2)	(0.2)	(1.4)	(1.7)	(1.8)
MKT	0.10	0.10	0.10	0.11	0.11	0.10	0.09	0.09	0.09
	(9.7)	(9.6)	(9.5)	(9.7)	(9.7)	(9.6)	(8.7)	(8.7)	(8.6)
MKT(-1)		0.01	0.01		0.01	0.02		0.00	0.00
		(1.4)	(1.4)		(1.3)	(1.4)		(0.2)	(0.3)
MKT(-2)			-0.01			0.00			0.00
			(-0.5)			(0.3)			(0.4)
Total MKT	0.10	0.11	0.11	0.11	0.12	0.12	0.09	0.10	0.10
SMB				0.03	0.03	0.03	0.03	0.03	0.03
				(2.1)	(1.7)	(1.8)	(2.0)	(1.6)	(1.7)
SMB(-1)					-0.02	-0.02		-0.01	-0.02
					(-1.3)	(-1.4)		(-0.9)	(-1.0)
SMB(-2)						-0.03			-0.02
						(-1.8)			(-1.6)
Total SMB				0.03	0.01	-0.02	0.03	0.01	-0.01
HML				0.07	0.07	0.07	0.05	0.05	0.05
				(4.2)	(3.9)	(4.0)	(3.0)	(2.8)	(2.9)
HML(-1)					0.01	0.01		0.00	0.00
					(0.8)	(0.7)		(-0.1)	(-0.1)
HML(-2)						0.00			0.00
						(-0.1)			(-0.2)
Total HML				0.07	0.08	0.08	0.05	0.05	0.04
UMD							-0.06	-0.05	-0.05
							(-5.2)	(-5.0)	(-4.9)
UMD(-1)								-0.03	-0.03
								(-2.7)	(-2.5)
UMD(-2)									-0.01
									(-0.7)
R2	14%	14%	14%	16%	16%	16%	20%	20%	20%

Table 4. Risk and Return Statistics (12/1974–10/2017)

This table reports the annualized returns, annualized volatility, annualized Sharpe Ratio and t-stat associated with the IGC premium calculated from the Bloomberg U.S. Corporate Index return series spliced with the excess return reported by Bloomberg starting in August 1988 for that same index. See body of paper for IGC premium computational details.

	IGC
Return	0.5
Volatility	3.5
Sharpe Ratio	0.17
t-stat	1.1

Table 5. Regression results (12/1974–10/2017)

This table reports the regression results associated with nine different regressions of the monthly returns of the IGC premium calculated from the Bloomberg U.S. Corporate Index return series spliced with the excess return reported by Bloomberg starting in August 1988 for that same index on various factor models with and without lagged variables. See body of paper for IGC premium computational details.

	I	II	III	IV	V	VI	VII	VIII	IX
Alpha	-0.0148	-0.0370	-0.0303	-0.0473	-0.0634	-0.0515	0.0039	0.0195	0.0416
	(-0.4)	(-0.9)	(-0.7)	(-1.2)	(-1.5)	(-1.2)	(0.1)	(0.5)	(1.0)
MKT	0.09	0.09	0.09	0.09	0.09	0.09	0.08	0.08	0.08
	(10.3)	(10.1)	(10.0)	(9.9)	(9.9)	(9.7)	(9.0)	(8.9)	(8.8)
MKT(-1)		0.04	0.04		0.03	0.03		0.02	0.02
		(3.9)	(3.9)		(3.2)	(3.2)		(1.9)	(1.9)
MKT(-2)			-0.01			-0.01			-0.01
			(-1.1)			(-1.0)			(-1.0)
Total MKT	0.09	0.13	0.12	0.09	0.12	0.11	0.08	0.10	0.09
SMB				0.06	0.05	0.05	0.06	0.05	0.05
				(4.2)	(3.5)	(3.5)	(4.7)	(3.9)	(3.8)
SMB(-1)					-0.01	-0.01		0.00	0.00
					(-0.8)	(-0.7)		(0.1)	(0.1)
SMB(-2)						-0.01			0.00
						(-0.8)			(-0.2)
Total SMB				0.06	0.04	0.03	0.06	0.05	0.05
HML				0.06	0.05	0.06	0.04	0.03	0.03
				(4.2)	(3.8)	(3.8)	(2.6)	(2.3)	(2.2)
HML(-1)					0.01	0.01		-0.01	-0.01
					(0.6)	(0.9)		(-0.6)	(-0.4)
HML(-2)						-0.02			-0.02
						(-1.1)			(-1.6)
Total HML				0.06	0.06	0.05	0.04	0.02	0.00
UMD							-0.06	-0.06	-0.06
							(-6.8)	(-6.4)	(-6.4)
UMD(-1)								-0.04	-0.04
								(-4.1)	(-4.1)
UMD(-2)									-0.01
									(-1.7)
R2	17%	19%	19%	21%	22%	22%	28%	30%	30%

Table 6. Ibbotson Long-Term Corporate Replication (1/1969–10/2017)

This table compares the risk and returns associated with the Ibbotson long-term corporate bond return series to that of "replicating" portfolios composed of stocks and government bonds. "Portfolio 1" allocates 14% to a market capitalization weighted portfolio of U.S. stocks and 86% to long-term government bonds. "Portfolio 2" allocates 9% to a market capitalization weighted portfolio of U.S. stocks, 5% to U.S. small-cap value stocks and 86% to long-term government bonds.

	Ibbotson LC	Portfolio 1	Portfolio 2
Average Return	0.70	0.72	0.74
Compound Return	8.31	8.50	8.74
Annualized SD	9.5	9.7	9.7
t-stat	2.8	2.8	3.0
Annualized Sharpe Ratio	0.39	0.41	0.43
Min. Return	-9.5	-10.8	-11.2
Max. Return	15.6	13.8	13.9
Max DD	-22.4	-17.7	-18.1
% Neg. Periods	37%	39%	39%
Skewness	0.41	0.31	0.28
Kurtosis	6.6	5.1	5.2

Table 7. Bloomberg U.S. Corporate Index Replication (1/1973–10/2017)

This table compares the risk and returns associated with the Bloomberg U.S. Corporate Bond Index return series to that of "replicating" portfolios composed of stocks and government bonds. "Portfolio 1" allocates 12% to a market capitalization weighted portfolio of U.S. stocks, 70% to Treasury bonds and 18% to long-term Treasury bonds. "Portfolio 2" allocates 8% to a market capitalization weighted portfolio of U.S. stocks, 4% to U.S. small-cap value stocks, 69% to Treasury bonds and 19% to long-term Treasury bonds.

	Bloomberg	Portfolio 1	Portfolio 2
Average Return	0.65	0.65	0.67
Compound Return	7.78	7.91	8.18
Annualized SD	7.0	5.9	5.9
t-stat	3.0	3.6	3.9
Annualized Sharpe Ratio	0.44	0.54	0.58
Min. Return	-8.1	-5.1	-5.3
Max. Return	12.9	9.9	10.0
Max DD	-19.3	-8.2	-8.6
% Neg. Periods	33%	34%	34%
Skewness	0.37	0.44	0.40
Kurtosis	8.2	5.5	5.6

Figure 1: Decade-by-Decade Sharpe Ratios

This figure graphs the decade-by-decade Sharpe Ratios associated with the equity, term and IGC premiums where the IGC premium is calculated using the Ibbotson long-term corporate bond return series. See body of paper for IGC premium computational details.

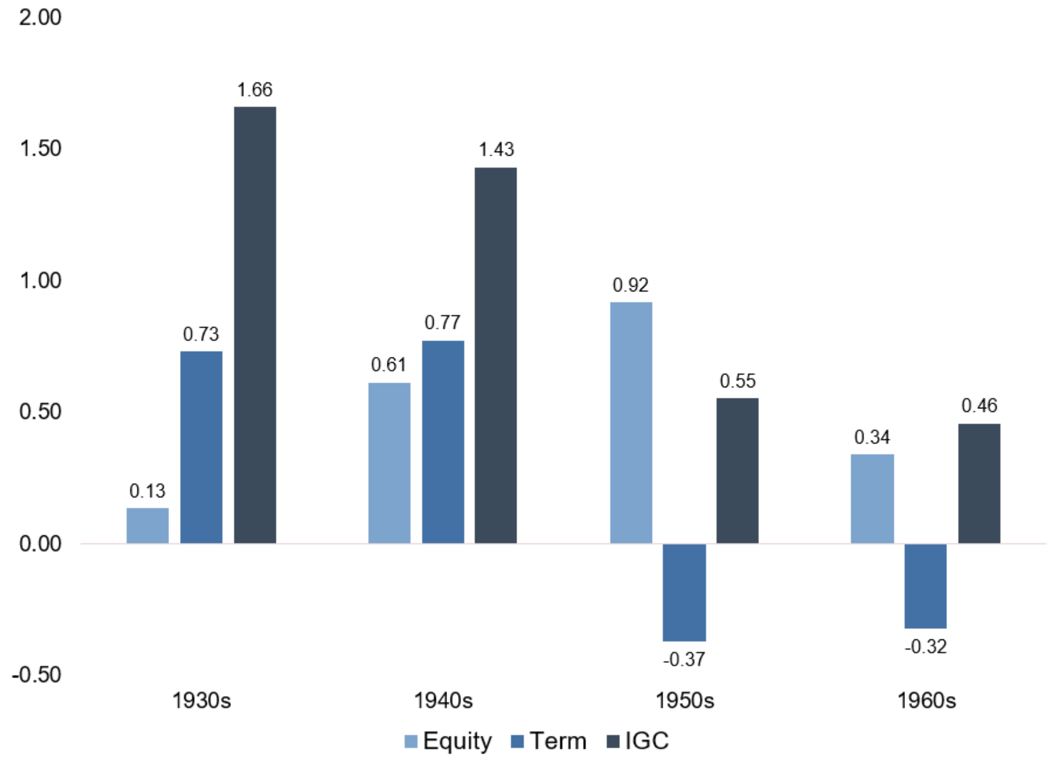

Figure 2: IGC Premium Annual Returns (1930–1968)

This figure graphs the year-by-year annual returns associated with the IGC premium calculated using the Ibbotson long-term corporate bond return series. See body of paper for IGC premium computational details.

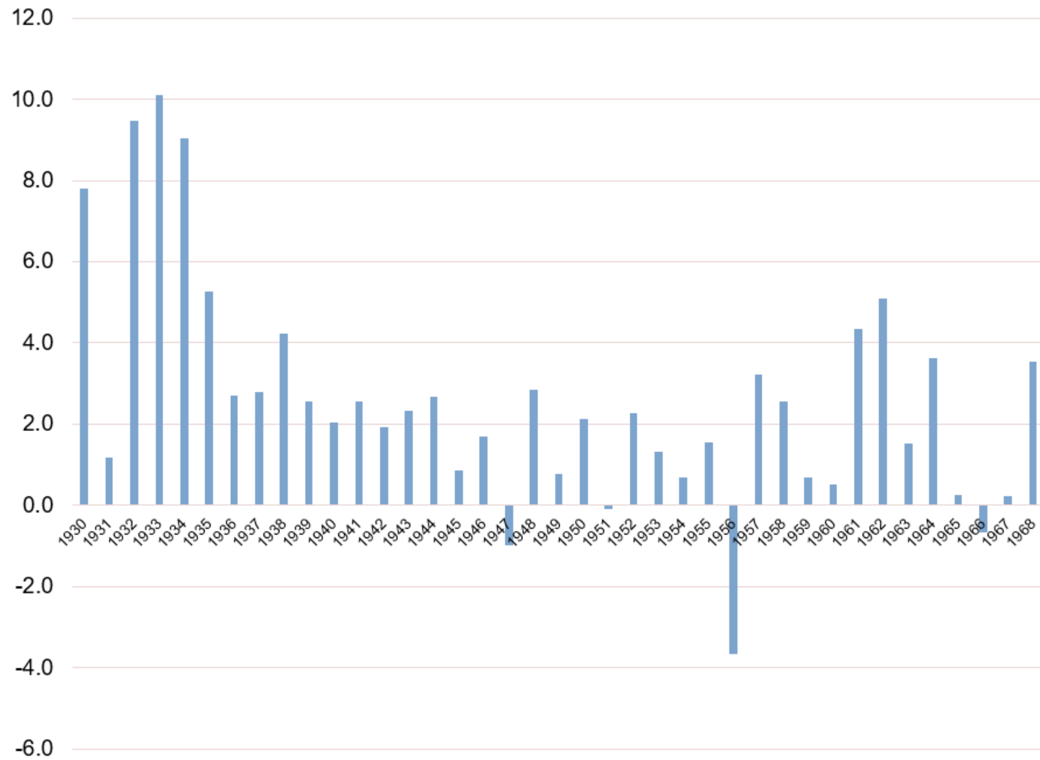

Figure 3: Investment-Grade Issuer Default Counts

This figure graphs the annual investment-grade default counts reported by Moody's in its 2016 study of the global corporate bond market default history.

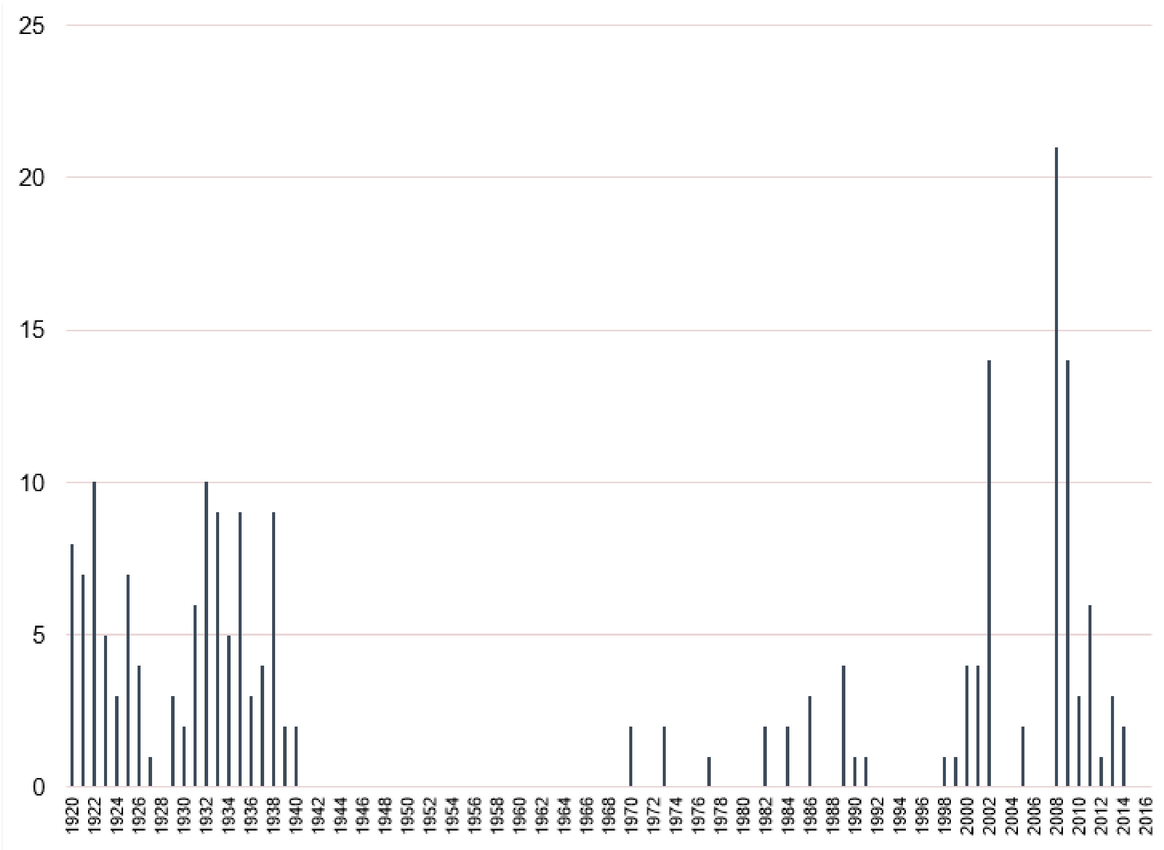

Figure 4: Speculative-Grade Issuer Default Counts

This figure graphs the annual speculative-grade default counts reported by Moody's in its 2016 study of the global corporate bond market default history.

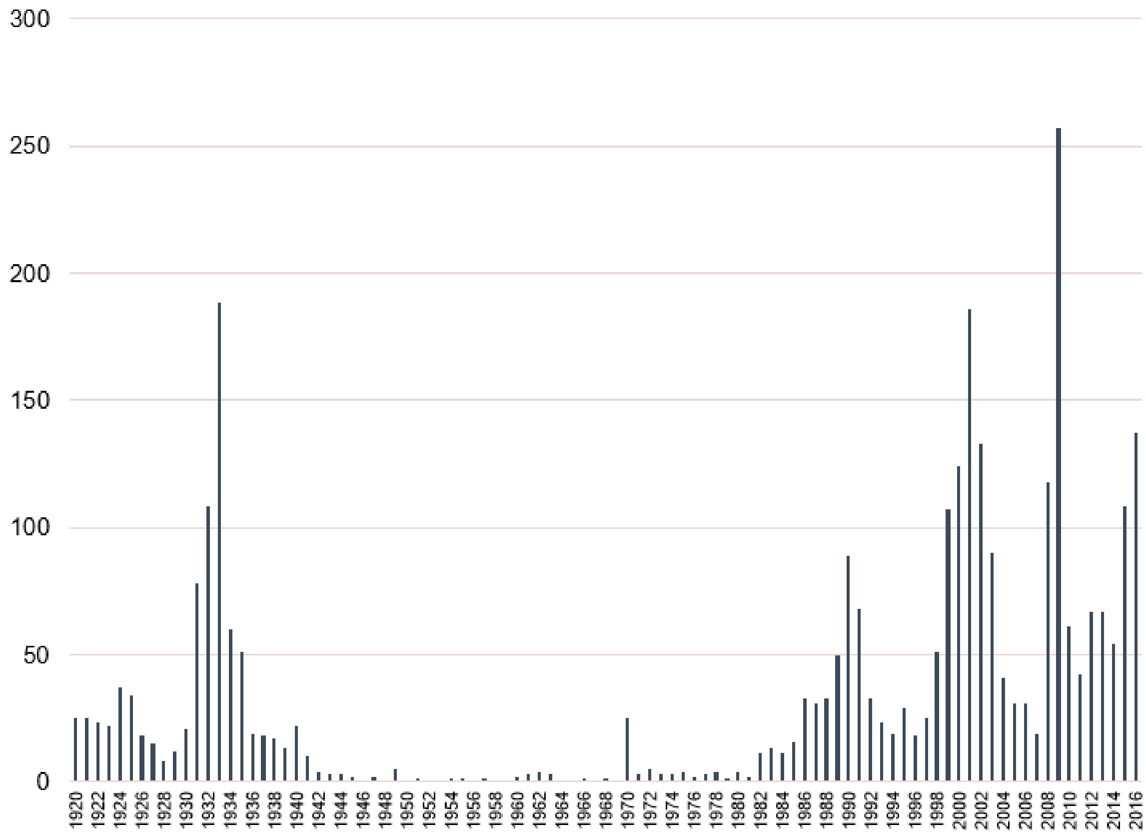

Figure 5: IGC Premium vs. IG Issuer Default Count (1930–1968)

This figure graphs a scatter plot of the annual IGC premium calculated using the Ibbotson long-term corporate bond return series against the Moody's reported annual default counts for global investment-grade corporate bonds. See body of paper for IGC premium computational details.

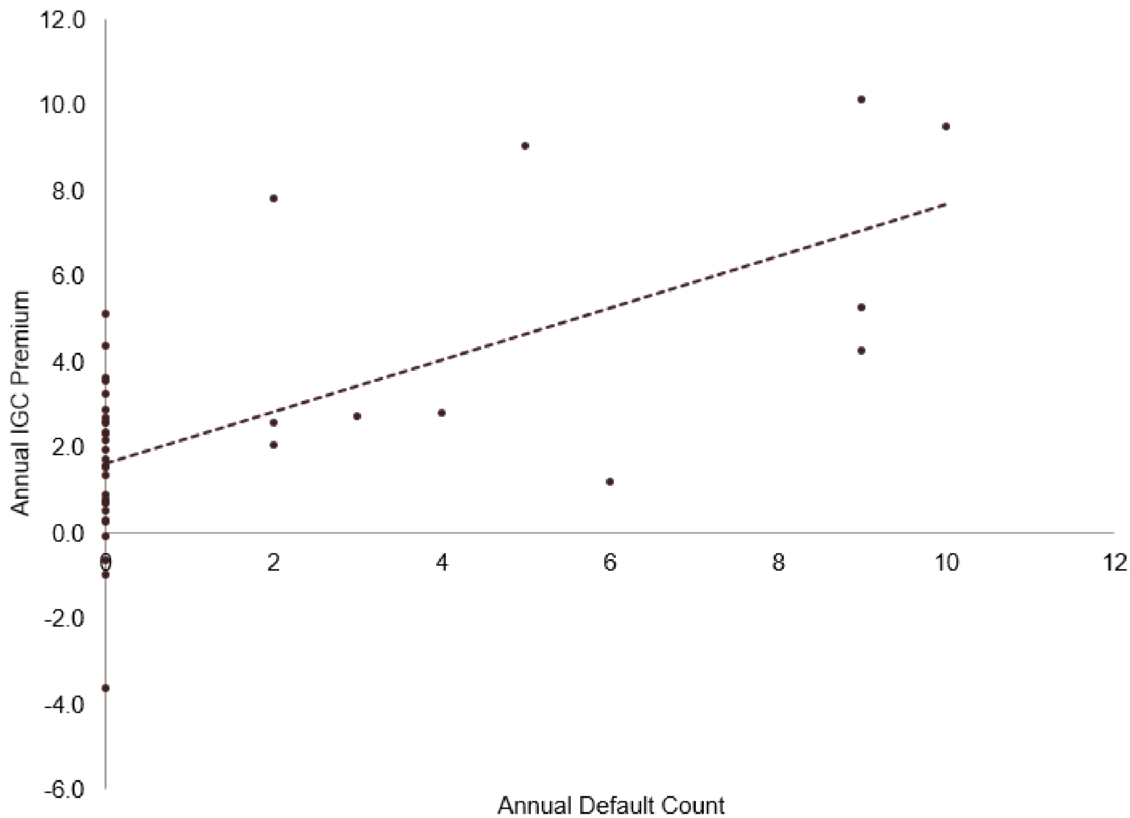

Figure 6: IGC Premium vs. IG Issuer Default Rate (1930–1968)

This figure graphs a scatter plot of the annual IGC premium calculated using the Ibbotson long-term corporate bond return series against the Moody's reported annual default rates for investment-grade corporate bonds. See body of paper for IGC premium computational details.

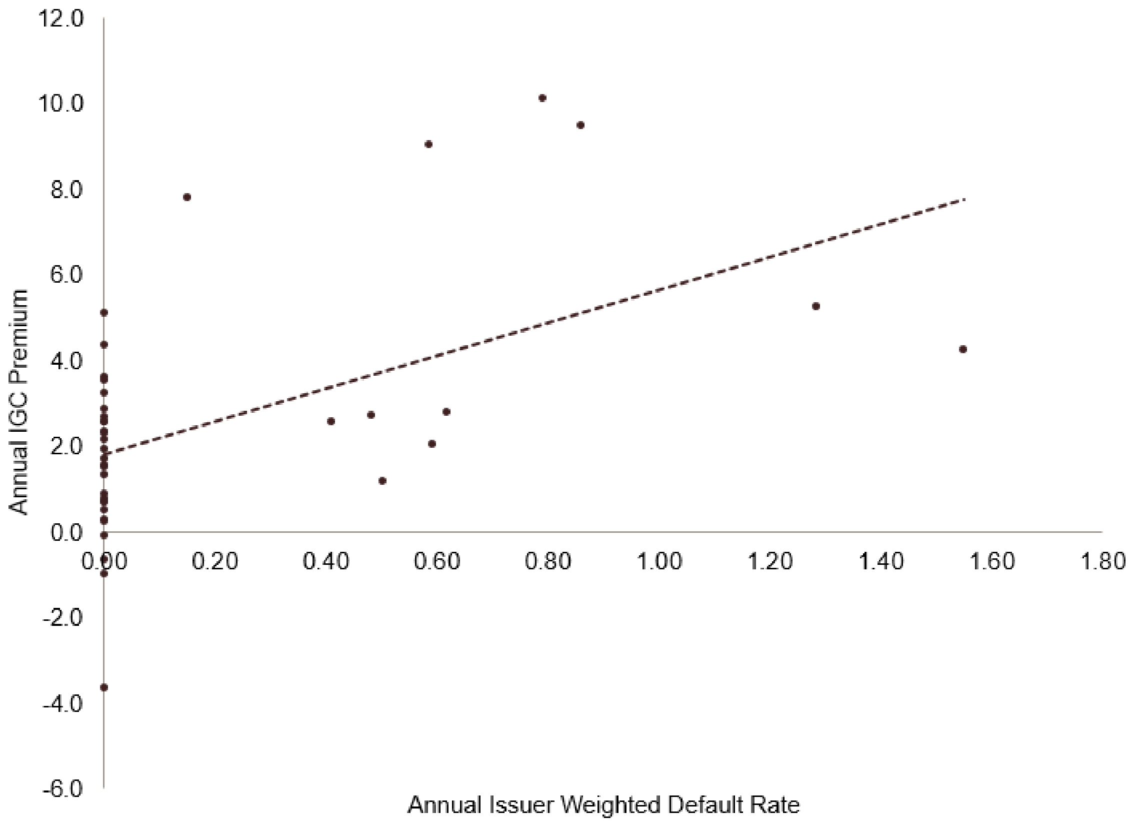

Figure 7: IGC Premium vs. IG Issuer Default Count (1969–2016)

This figure graphs a scatter plot of the annual IGC premium calculated using the Ibbotson long-term corporate bond return series spliced with the excess returns reported by Bloomberg for the Bloomberg U.S. Corporate Index against the Moody's reported annual default counts for global investment-grade corporate bonds. See body of paper for IGC premium computational details.

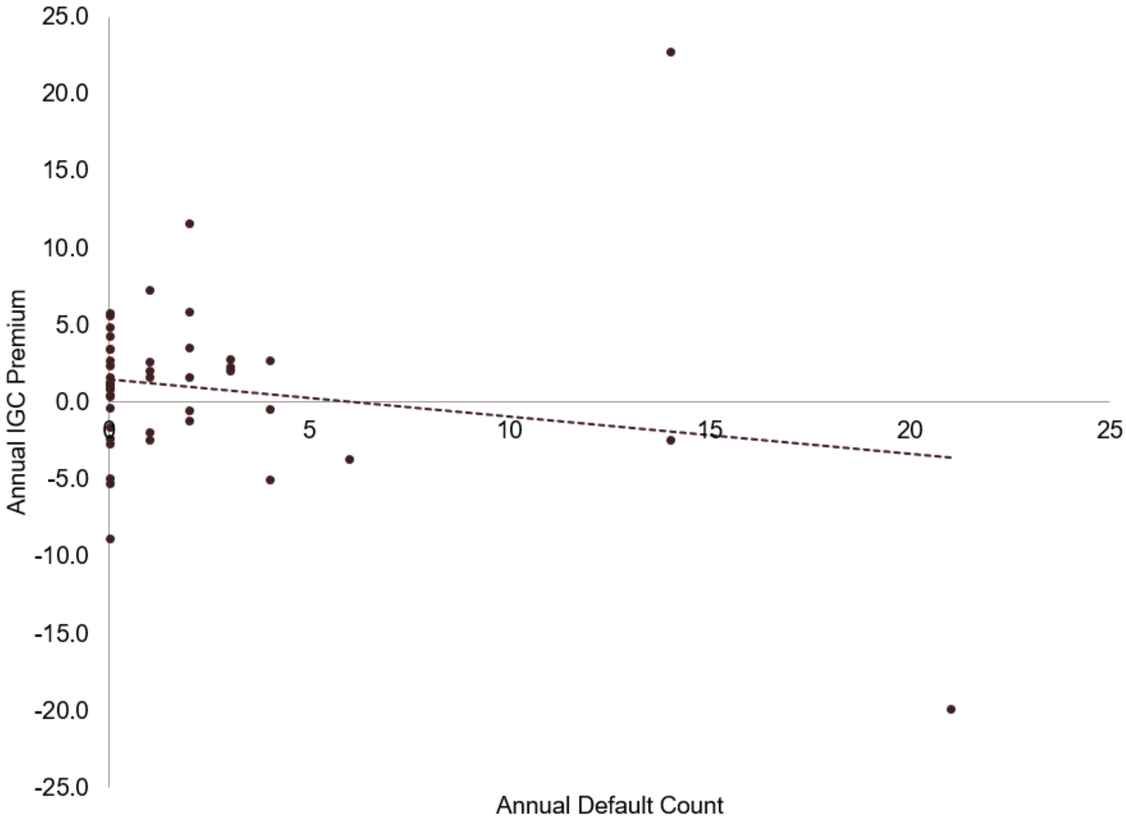

Figure 8: IGC Premium vs. IG Issuer Default Rate (1969–2016)

This figure graphs a scatter plot of the annual IGC premium calculated using the Ibbotson long-term corporate bond return series spliced with the excess returns reported by Bloomberg for the Bloomberg U.S. Corporate Index against the Moody's reported annual default rate for investment-grade corporate bonds. See body of paper for IGC premium computational details.

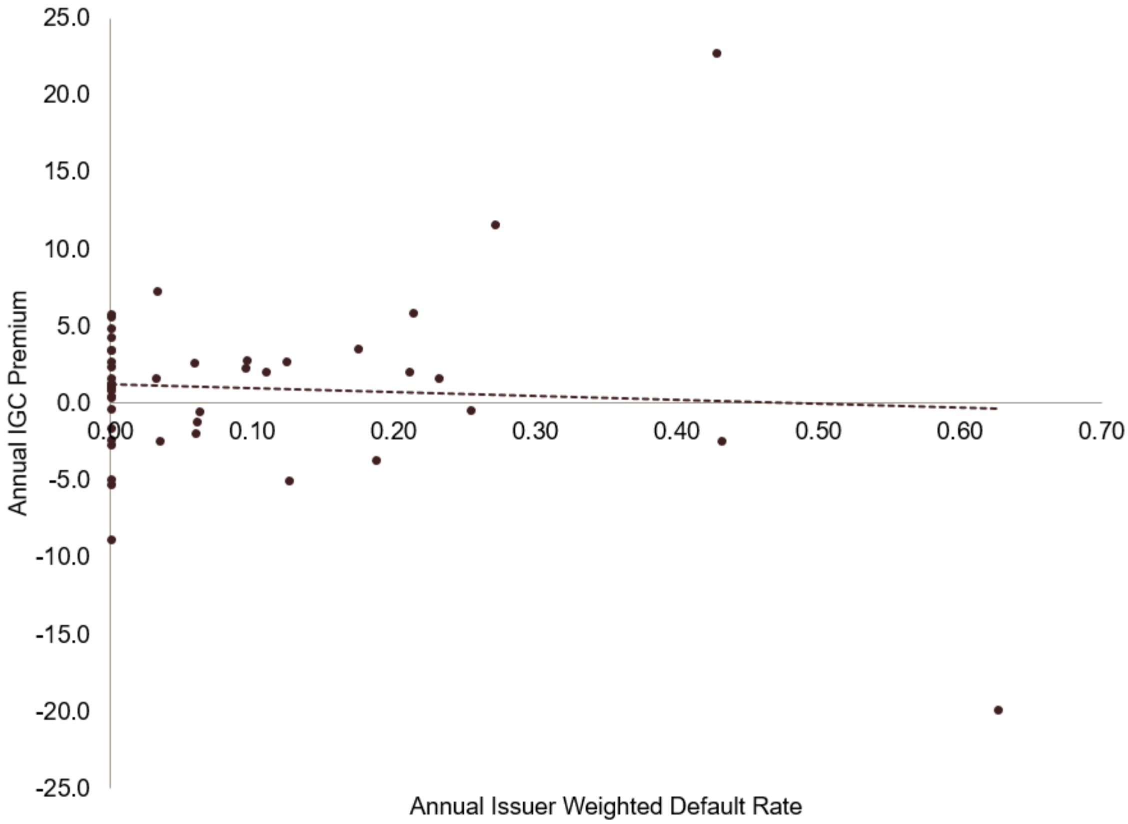

Figure 9: IGC Premium vs. IG Issuer Default Count (1975–2016)

This figure graphs a scatter plot of the annual IGC premium calculated using the Bloomberg U.S. Corporate Index spliced with the excess returns reported by Bloomberg for that same index against the Moody's reported annual default counts for global investment-grade corporate bonds. See body of paper for IGC premium computational details.

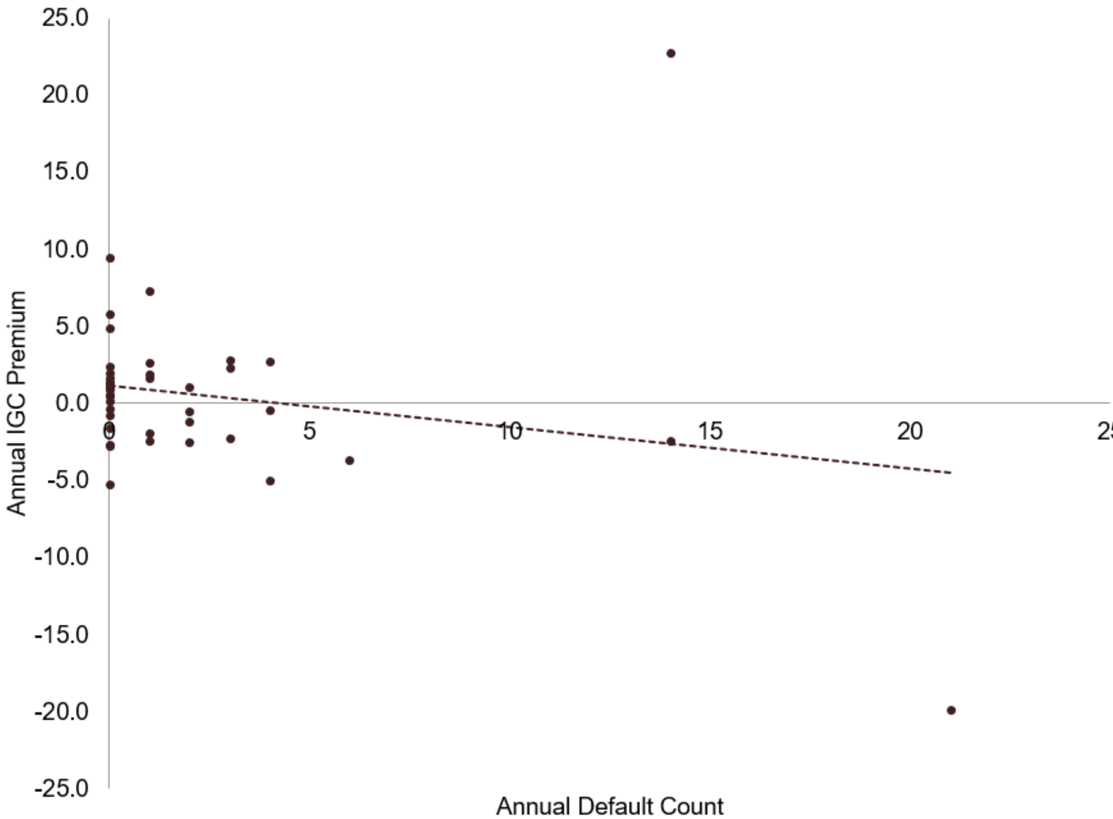

Figure 10: IGC Premium vs. IG Issuer Default Rate (1975–2016)

This figure graphs a scatter plot of the annual IGC premium calculated using the Bloomberg U.S. Corporate Index spliced with the excess returns reported by Bloomberg for that same index against the Moody's reported annual default rates for investment-grade corporate bonds. See body of paper for IGC premium computational details.