APPENDIX 12

Informing a Savings Target using Revenue Volatility

A report by The Pew Charitable Trusts to the Wyoming Joint Revenue Committee

September 11, 2015

Executive Summary

At the request of the Joint Revenue Committee, The Pew Charitable Trusts has worked collaboratively with the Wyoming Legislative Service Office (LSO) to assess the volatility of the state's revenue portfolio in order to inform the Legislature's determination of an evidence-based savings target for the state's budget reserves, including the Legislative Stabilization Reserve Account (LSRA) and other funds. The study presented in this technical report is an adaptation of the methodology used by the State of Minnesota to set their annual savings target. Pew examined Wyoming's revenue portfolio across six major revenue streams: sales and use taxes, severance taxes, investment income, mineral royalties, K-12 property taxes, and "other" revenues. All data used in this analysis were provided by the Wyoming Legislative Services Office, and data considerations are detailed in the appendix of the report.

The Wyoming model presented in this analysis does not offer one, fixed savings target recommendation, but instead requires policymakers to answer two questions related to the level of budgetary risk they wish to protect against with the fund in order to arrive at a recommendation. This choice is similar to deciding what level of insurance coverage the state desires, with higher levels of protection requiring a higher premium and vice versa. In order to choose the ideal savings target for Wyoming's reserves, policymakers must first decide how much budgetary risk they wish to offset with the fund's balance ("level of coverage"), and second, how many years' worth of revenue downturns they wish to protect against ("downturn duration").

- **Level of coverage** is characterized by a level of protection against revenue downturns in the form of a percentage. Similar to an insurance policy, a more comprehensive policy – for example, one that covers all but 10 percent of the potential revenue downturns (i.e., a 90 percent level of protection) – will "cost" more in the form of a higher recommended level of savings than a less comprehensive policy with a lower level of protection, such as all but 20 or 50 percent of revenue downturns (i.e., 80 percent or 50 percent levels of protection, respectively).
- **Downturn duration** is the number of years a revenue downturn's impact carries forward through Wyoming's budget. This analysis draws from the state's historical experience to estimate the persistence of a single-year revenue loss over time – namely, how long a downturn in one year carries over into subsequent years. However, at high levels of coverage, the savings target suggested by the model would be adequate to smooth over either a large, single-year revenue loss or multiple years of relatively smaller downturns.

The estimates presented below reflect a range of options for the desired level of coverage that can be provided by the state's reserves as well as the downturn duration deemed appropriate. Both of these parameters can be adjusted to reflect policymakers' preferences.

Table 1
Wyoming Reserve Funds
Recommended Savings Targets at Various Levels of Coverage and Downturn Durations

	Level of Coverage for Revenue Downturns					
Downturn Duration (in years)	**50%**		**80%**		**90%**	
1	13.0%	$591.7	24.4%	$1,113.8	31.4%	$1,435.7
2	20.3%	$925.3	38.1%	$1,741.8	49.2%	$2,245.3
5	28.0%	$1,279.3	52.7%	$2,408.2	68.0%	$3,104.3
10	29.6%	$1,352.3	55.7%	$2,545.4	71.9%	$3,281.2

Note: Values reported are percentages of annual revenues from the designated portfolio. Dollar amounts reported are in millions. Level of coverage reflects the amount that would be needed to insulate Wyoming's budget from a set percentage of possible downturns, with higher values indicating greater levels of protection.

Given Wyoming's exceptionally high levels of revenue volatility and unique tax portfolio, as well as challenges with data availability, adapting Minnesota's methodology required a number of revisions in order to produce appropriate estimates of Wyoming's revenue volatility and potential savings targets. The Pew project team consulted extensively with economist Matthew Schoeppner of Minnesota Management and Budget, former Minnesota state economist Dr. Tom Stinson of the University of Minnesota, and Don Richards of the Wyoming Legislative Services Office (Wyoming LSO) in order to ensure that these adaptations maintained the analytical rigor of the source methodology.

This paper reviews Wyoming's full revenue portfolio and goes through this adapted Minnesota model with a step-by-step description of how the estimates in Table 1 were generated.

Wyoming's Revenue Portfolio

Previous examinations of Wyoming's revenues have provided valuable insights into the volatility in severance collections, energy markets, and the state's revenue experience. However, under the direction of the Revenue Committee and Legislative Service Office, we examined an "expanded" revenue portfolio that evaluates revenues across the state's traditional funds (General Fund and Budget Reserve Account) in addition to funding for K-12 education (operations funded from the School Foundation Program Account and capital construction funded from the School Capital Construction Account), and total severance tax and federal mineral royalty (and coal lease bonus) streams. This revenue portfolio is displayed in the figure below (Figure 1).

Figure 1
Wyoming's Revenue Portfolio
Shares of total revenues, 1970-2014

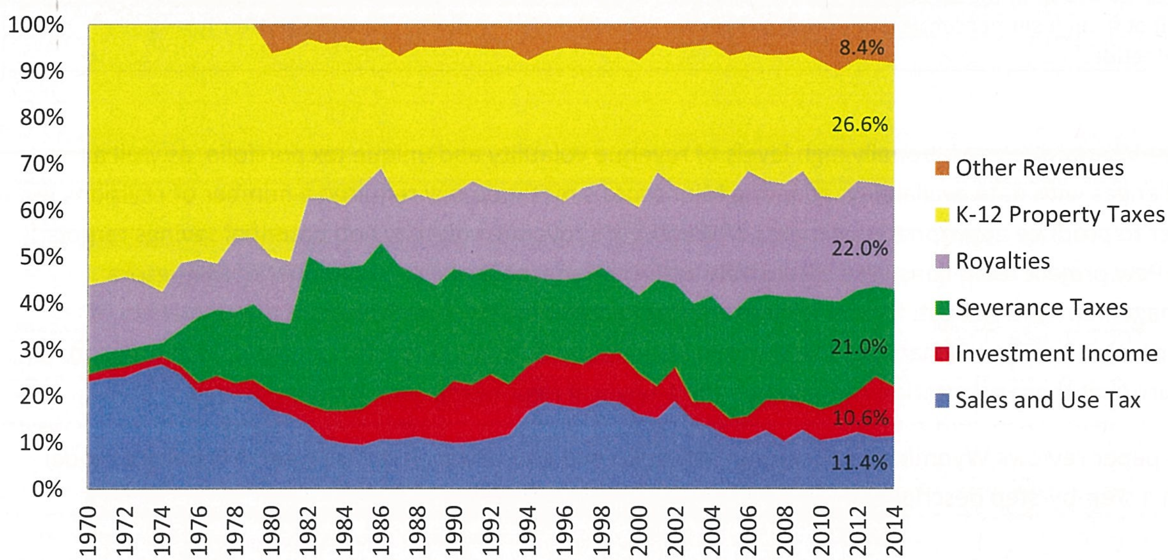

For the period between 1970 and 1982, the portfolio is dominated by the contributions of K-12 funding property taxes. However, this share is reduced considerably between 1980 and 1984, largely due to the introduction of the Coal Lease Bonus Payment (CLBP) and "other" revenue streams.[1] In addition, there is a sharp increase in the proportion of revenues attributable to severance tax collections in 1981. While the relative share of each stream largely stabilized by the 1990s, significant fluctuations continued to reshape the portfolio. For example, the relative importance of the state's sales and use tax increased in 1994 and remained at elevated levels through the early 2000s when it returned to its current share of 11.4 percent.

---

[1] The "other" revenues stream includes state revenues attributable to smaller revenue streams that do not fall within the confines of the major revenue streams. Examples of revenues included in this stream are cigarette taxes, franchise taxes, excise tax on wind generated electricity, charges for sales and service, penalties and interest assessed on unpaid taxes, state royalties, and many more.

Although the composition of the portfolio in any given year does not offer a full portrait of the volatility in each respective stream, the sheer amount of variation in each stream emphasizes the need to examine volatility more closely and over time.

Why Pew Chose to Adapt the Minnesota Approach for Wyoming

In a 50-state analysis of state tax volatility published as part of Pew's Fiscal 50: State Trends and Analysis Interactive, Pew examined a 19-year period and presented a "volatility score" – namely, a calculation of a standard deviation in revenue changes. Pew's Fiscal 50 interactive tool includes key fiscal, economic, and demographic trends in the 50 states and allows users to analyze their impact on states' fiscal health; this provides an informative examination of how a state's revenue portfolio fluctuates over time. However, the ability to generate a savings target using these estimates is limited. This is because a simple standard deviation treats all the observed data identically when, as evident in the figure displaying Wyoming's revenue portfolio over time, this is not the case. The Minnesota approach offers a more comprehensive assessment with the following benefits:

- Differentiation between revenue streams
- Accounts for the composition of the revenue portfolio
- Recognition of the covariation between revenue streams
- Adjustment to account for the passage of time

For these reasons, Pew has adapted the Minnesota model to the state of Wyoming and its highly volatile, mineral-centric revenue portfolio in order to model optimal savings target options for the Legislative Stabilization Reserve Account and other state reserves.

Step 1: Identifying the Cyclical Trend

The first step in the Minnesota model is to **separate out the cyclical trend** from each individual revenue series. This allows for the removal of any volatility that would come from the regular business cycle in our assessment of the state's revenue streams. To accomplish this task, we examine the log returns of each revenue series and calculate the trend in each series using a Hodrick-Prescott (HP) filter.[2] The

[2] Transforming the series from raw, nominal dollar figures to logarithmic values allows for year-over-year changes to approximate percentage changes and mitigates the distortions caused by balances changing significantly from year-to-year, which is especially problematic when revenues drop to low values or even zero as is the case for Wyoming's Coal Lease Bonus Payments. For example, if a revenue stream only had $5 million one year and $125 million the next, then that would report a 12,400 percent increase. For log-transformed data, the same increase would report a more manageable 382.8 percent increase. Note that the increase is still sizeable, but much more

figure below (Figure 2) displays the observed returns for Wyoming's sales and use tax from 1971 through 2014 as compared to the cyclical trend calculated using the HP filter. Figures for the other revenue streams are available in the technical appendix.

Figure 2
Wyoming's Sales and Use Tax
Observed returns and trend, 1971-2014

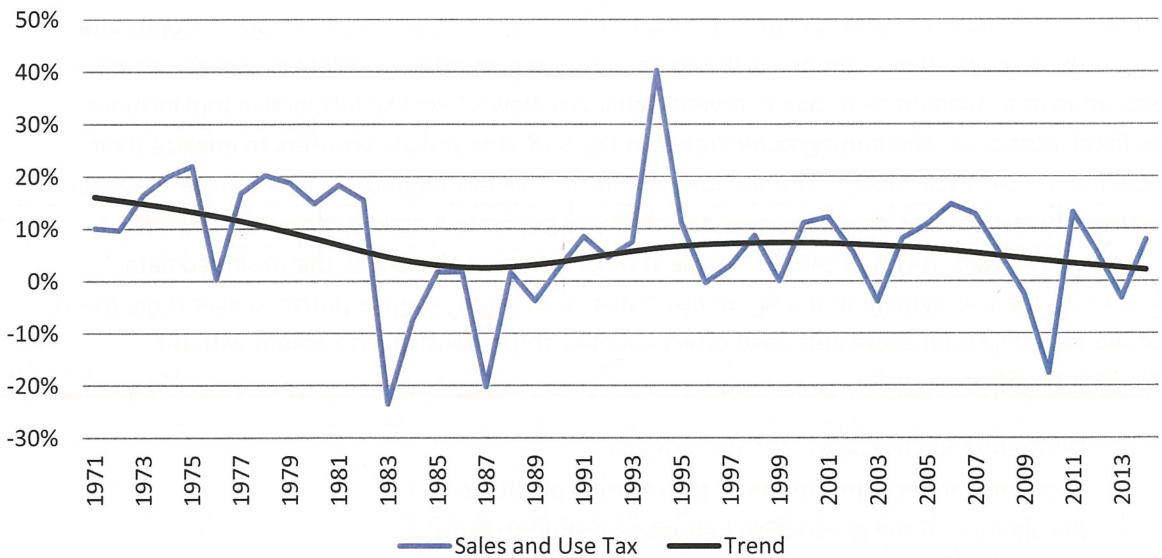

The first point of note in Figure 2 is the volatility in the observed statewide sales and use taxes deposited into the General Fund. Changes appear to have been largely predictable, with growth ranging from 10-20 percent throughout the 1970s and early 1980s, but this changed abruptly in 1983, when sales and use dropped from a pattern of moderate growth to contraction. Throughout the 1980s, sales and use taxes appear to shrink considerably, with annual collections observing about 20 percent declines in both 1983 and 1987 and little-to-zero growth in the interim. While the 1990s marked a return to growth in sales tax revenues, there is a notable spike in 1994 which corresponds with an increase in the tax rate. Perhaps more important than the nuances of the observed changes in revenues is the HP trend line separating the data into the long-run cyclical trend line and the year-to-year deviations from that trend. . The trend demonstrates the contributions the business cycle made to the changes in sales and use revenues for the state. The sharp drop observed in 1983 appears in the midst of a decline in the business cycle, suggesting that a correction of some sort was likely to have occurred.

For our purposes, we are less interested in the particulars of this revenue stream; rather, we use these data to observe and calculate the deviations from trend, illustrated by the red bars in the figure below (Figure 3).

---

manageable within the context of estimating volatility. Additionally, the use of log returns is a well-established practice within economics and the return data approximate year-over-year percentage changes in most cases.

Figure 3
Wyoming's Sales and Use Tax
Observed returns and trend, 1971-2014

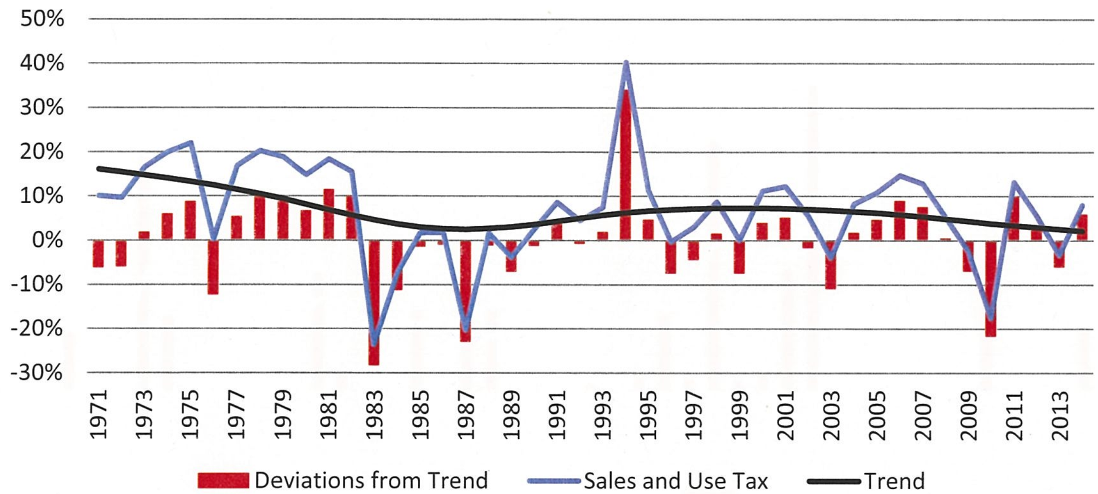

Step 2: Calculating Time-Varying Cyclical Volatility

The second step is to use the annual deviations from trend calculated in the previous step to generate a measure of volatility for each revenue stream. We do this by **calculating "time-varying" cyclical volatility**, called as such because the variance for each yearly observation is based on variance in the years preceding that observation.[3] The figure below (Figure 4) presents the time-varying standard deviation that is produced by the model for the sales and use taxes for the state. To place this standard deviation into context, it is plotted against the absolute value of the deviations-from-trend for the series.[4]

[3] An Integrated Generalized Auto-Regressive Conditional Heteroskedasticity (IGARCH) model was employed to produce these estimates. This is a widely used model of volatility adapted from finance and investment banking that allows the variance in a series to change over-time. This is useful in this context as the volatility in state revenues is not constant over time, as illustrated in Figure 2, where the variance in the state's sales and use taxes is considerably less in the 1990s and early 2000s than the more volatile 1980s and 2010s.

[4] The deviations-from-trend are presented as absolute values for graphical reference only. They are **NOT** transformed during or for the estimation of the IGARCH.

Figure 4
Characteristics of Wyoming's Sales and Use Tax
Deviations and conditional standard deviation, 1971-2014

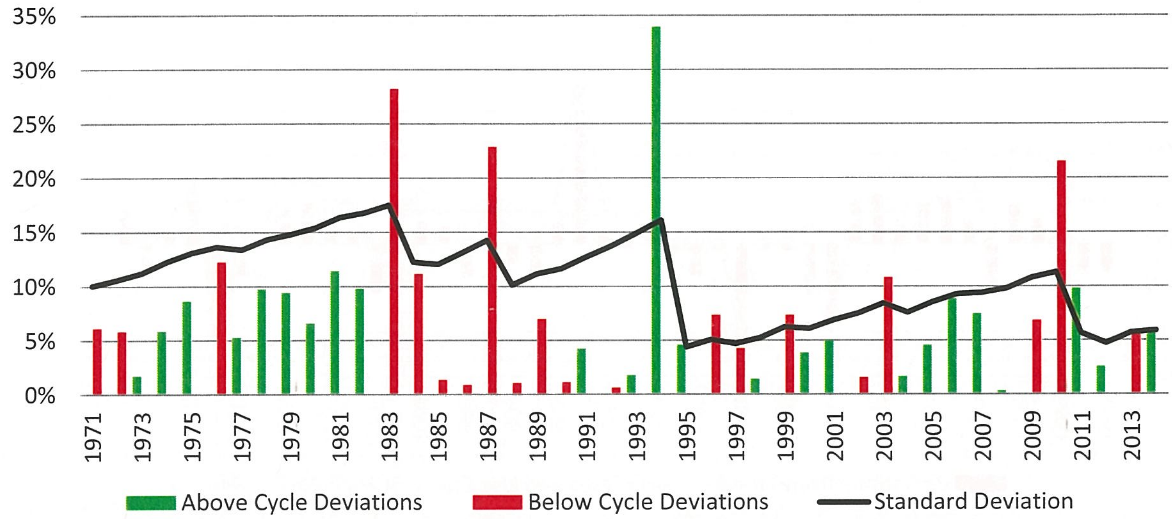

As the figure illustrates, sales and use tax exhibit the largest deviations from trend in 1983, 1987, 1994, and 2010, corresponding to the previous figures (Figures 2 and 3). Further, the standard deviation appears to track the data well, updating in the years following significant spikes away from the trend. Lastly, it is appropriate to interpret this time-varying standard deviation as the volatility of the state's sales and use tax revenues from 1971 through 2014.

Step 3: Calculating Conditional Correlations between Revenue Streams

Having determined time-varying cyclical volatility for each stream, the next step of the process requires an examination of how each tax stream moves in relation to the other tax streams. Understanding these relationships helps to assess how the volatility of the overall portfolio may expand or contract over time. If two revenue streams covary highly, much like when two items in a stock portfolio follow the same cycle, the volatility for the entire portfolio increases. Conversely, when revenue streams do not move together, risk is mitigated and the portfolio volatility decreases. In order to measure the covariance between streams, we generate a **dynamic conditional correlation** for each possible pairing of revenue series.[5,6] The figures below (Figure 5) display the time-varying standard deviations and correlations between sales and use taxes and the other major revenue streams.

---

[5] This process, as detailed by Stinson and Schoeppner for Minnesota was developed by Engel (2002). First, the residuals from the IGARCH model for each stream were standardized with their respective conditional standard deviations. We then calculated the product of the standardized residuals for both series, and then estimated another IGARCH on the product series. We predicted the conditional variance of the product series and then calculated the conditional correlation.

[6] Note that we do not correlate every possible pairing. It is appropriate to estimate separate portfolios for the investment, severance, and mineral royalties revenues as each major stream is comprised of several sub-streams

Figure 5
Sales and Use Tax Correlations
Compared to other major revenue streams, 1972-2014

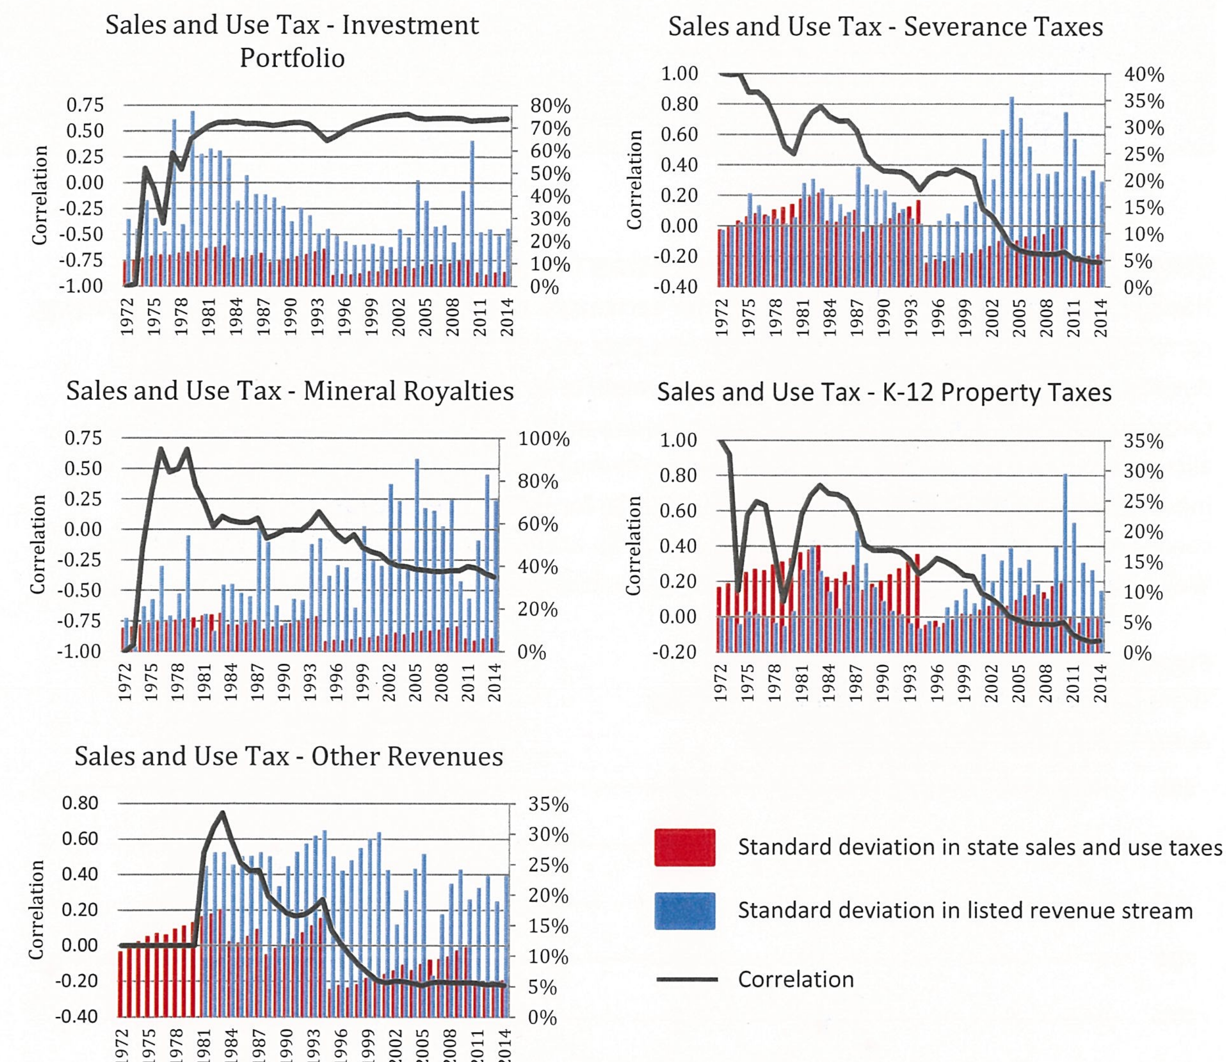

As the figures illustrate, the correlation between the volatilities in the state's sales and use tax and the other revenue streams are highly variant from 1970 through 2014. While the correlation between volatility in the sales tax and investment income reaches stability in the late 1990s, many of the correlations with other streams continue to vary up to present day. The exact correlations between sales and use taxes and the other revenue streams for fiscal year 2014 are presented in the table below (Table 2).

(i.e., oil, coal, natural gas, and other minerals). We estimate these sub-portfolios following the steps of the MN methodology and then insert the sub-portfolios as major streams into the analysis at this point.

Table 2
Fiscal Year 2014 Correlations between Major Revenue Streams

	Sales Tax	Investments	Severance	Royalties	Property	Other
Sales Tax	1.00	0.62	-0.24	-0.39	-0.13	-0.22
Investments	---	1.00	0.05	-0.30	0.19	-0.22
Severance	---	---	1.00	0.44	0.84	0.28
Royalties	---	---	---	1.00	0.26	0.24
Property	---	---	---	---	1.00	0.33
Other	---	---	---	---	---	1.00

Step 4: Time-Varying Standard Deviation for All Revenue Streams

Having calculated the time-varying volatility for each major revenue stream as well as the time-varying correlations between different streams, we were then able to assess the time-varying standard deviation of *all* revenue streams – a measure of volatility for the entire revenue portfolio. This was calculated using a **portfolio choice model** modified to account for variation over time.[7] This model allows us to estimate a standard deviation for the overall revenue portfolio for each fiscal year. The model incorporates the time-varying cyclical volatility for each respective tax stream (Step 2), conditional correlations between tax streams (Step 3), and weights for each respective tax stream. Weights were calculated by looking at each revenue stream's share of the total for each fiscal year.

Figure 6
Wyoming's Full Revenue Portfolio Volatility
Across all revenue streams, 1971-2014

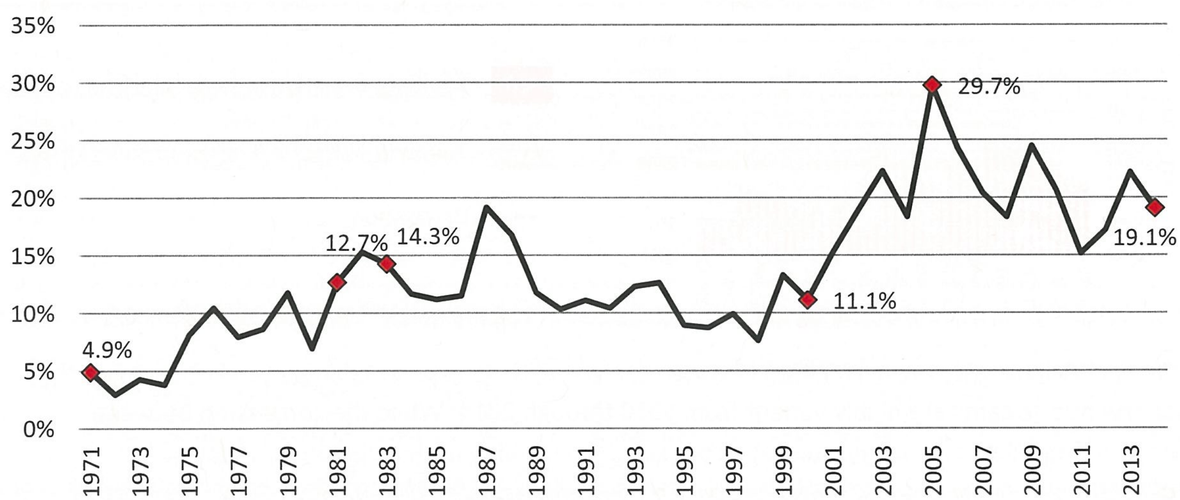

[7] The adaptation for time-variance was developed by Stinson and Schoeppner (2014).

The portfolio standard deviation exhibits considerable growth over the 40-year period from 1971 through 2014. The initial portfolio volatility in 1971 is identified at 4.9 percent of portfolio revenues; however, by the start of the 1980s, the volatility had increased almost three-fold, reaching 12.7 percent in 1981 when the "other" revenues and coal lease bonus payments were added to the portfolio. In 1983, the year Wyoming's Revenue Volatility and Savings Analysis identified as the start of a recession, the portfolio volatility was 14.3 percent. Volatility decreases over the course of the identified 17-year recession, reaching 11.1 percent by 2000. However, volatility in the portfolio increases dramatically driven by shifts in the mineral royalties stream between 2000 and 2005. In 2005, the portfolio reaches its peak volatility of 29.7 percent. Over the course of the Great Recession and years following the portfolio's volatility declined, reaching an estimated standard deviation of 19.1 percent of portfolio revenues for FY2014.

Step 5: Identifying an Optimal Fund Target for Wyoming's Reserves

The final step in this analysis was to translate the time-varying standard deviation presented by the modified portfolio model into a workable funding target for the LSRA and other state reserves. This required multiplying the figure by two factors – first, the level of risk the state intends to guard against, and, second, the period of time the balance should offset.

For example, in Stinson and Schoeppner's original model used for Minnesota, the authors assumed a 95 percent level of coverage. This meant that Minnesota's fund would be sufficiently large to cover 9 out of 10 possible downturns the state might encounter.[8] This level of coverage requires that the portfolio volatility is multiplied by a "critical value" of 1.65.[9] This was then multiplied by the square root of 2 in order to account for Minnesota operating on a biennial budget. This produced an estimate of 9.8 percent of annual revenues.

For Wyoming, the results of the full portfolio indicate a standard deviation for FY 2014 of 19.1 percent of annual revenues. Transforming these results to an appropriate fund target depends on what level of risk and what time horizon policymakers believe to be ideal. The table below (Table 3) presents a full range of options given different assumptions about levels of coverage and downturn durations.

---

[8] The level of coverage critical value draws upon a normal distribution of all deviations from the revenue trend. As such, the 95% level of coverage selected by Minnesota provides the conversion to offer protection against 95% of all estimated shocks to the state's revenues; however, half of these shocks are estimated surpluses. As such, while a 95% level of coverage corresponds to 19 out of 20 deviations (10 surpluses and 9 downturns), the protection only insulates from 9 out of 10 possible downturns.

[9] For more specific information regarding the identification of "critical values" for different levels of coverage, please reference Appendix B.

Table 3
Possible Savings Targets for Wyoming's Reserve Funds

	Level of Coverage for Revenue Downturns					
Downturn Duration (in years)	**50%**		**80%**		**90%**	
1	13.0%	$591.7	24.4%	$1,113.8	31.4%	$1,435.7
2	20.3%	$925.3	38.1%	$1,741.8	49.2%	$2,245.3
5	28.0%	$1,279.3	52.7%	$2,408.2	68.0%	$3,104.3
10	29.6%	$1,352.3	55.7%	$2,545.4	71.9%	$3,281.2

Note: Values reported are percentages of annual revenues from the designated portfolio. Dollar amounts reported are in millions. Level of coverage reflects the amount that would be needed to insulate Wyoming's budget from a set percentage of possible downturns, with higher values indicating greater levels of protection.

To provide a better understanding of how the values in the table were calculated, we present the savings targets for the state's reserves under two scenarios. The first scenario assumes a 50 percent level of coverage, which would offer complete protection against half of all possible downturns. A single year downturn is estimated and then simulated with a 10-year downturn duration using an estimate for how long shocks to the state's revenues persist in the portfolio.

Figure 6
Downturn Scenario 1
50% level of coverage, 10-year window

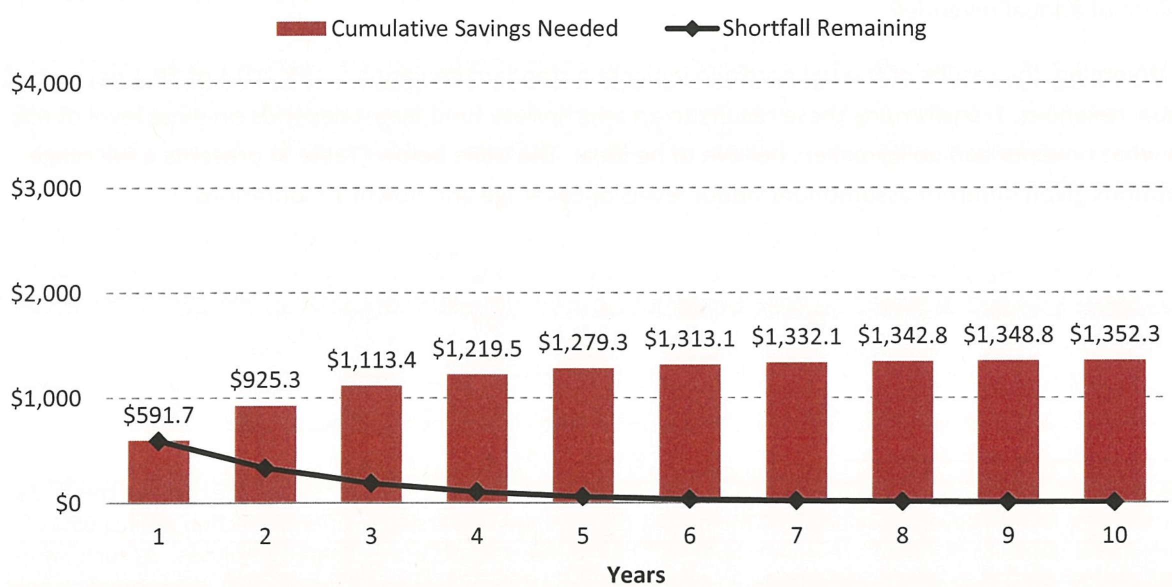

From this simulation we find the initial downturn would total $591.7 million in the first year. As time passes, the impact of the downturn diminishes rather quickly; five or six years after the initial downturn, the impact reduces to zero. However, we note the cumulative fiscal impact of the downturn is considerable, with total impact of the downturn eventually totaling more than $1.3 billion. As such, if the reserves were desired to carry this level of coverage, an estimated savings target for a 10-year downturn duration would be roughly $1.4 billion.

Figure 7
Downturn Scenario 2
90% level of coverage, 10-year window

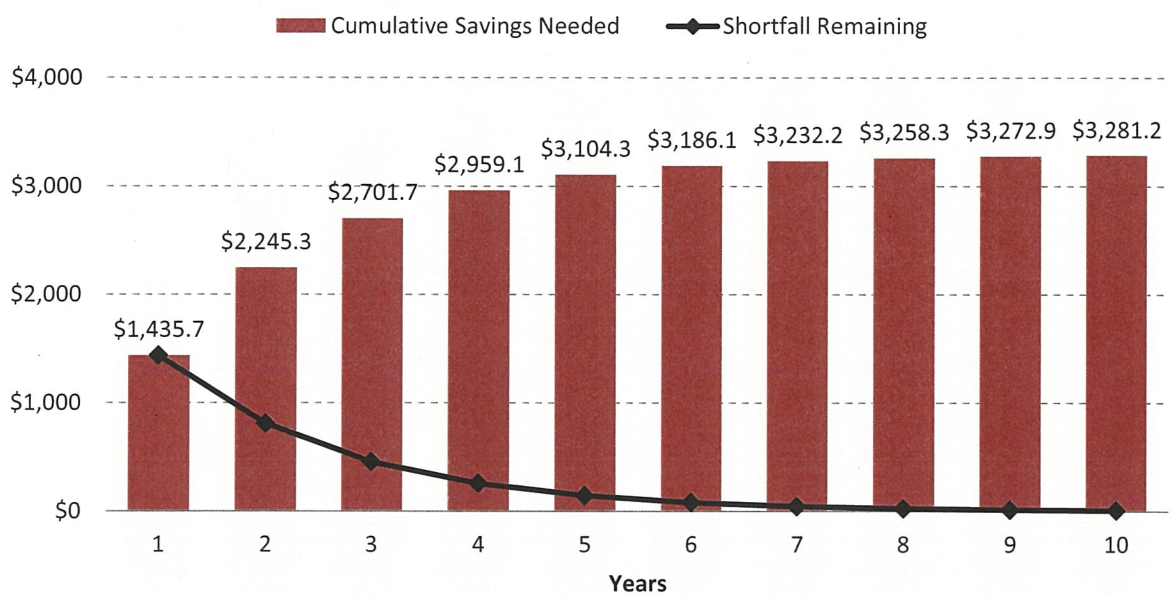

If we modify our assumed level of coverage provided by the reserves to match the protection offered in Minnesota, a 90 percent level of coverage, the estimated savings targets are increased considerably. Much like how a larger insurance policy requires a bigger monthly premium, if the state desires to expand their protection to 9 out of every 10 estimated downturns, the targets for the state's reserves increase accordingly. Under this simulation, the initial downturn estimate is more than $1.4 billion. But, similar to our first simulation, the full impact of the downturn is measured by how its effects persist in the revenue portfolio over the downturn duration. While the rate of decay is the same as the prior simulation, the estimated impact is considerably larger due to the initial downturn estimate. Assuming a 90 percent level of coverage and a 10-year downturn duration, the savings target would be nearly $3.3 billion.

Conclusion

This analysis of the state's revenue portfolio illustrates the high levels of revenue volatility that can be attributed to Wyoming's heavy reliance on mineral extraction. Through the adaptation of the Minnesota methodology, Pew found that the standard deviation in annual portfolio returns (a measure of volatility) equals 19.1 percent of the annual revenues. Analysis of the persistence of revenue downturns in

Wyoming's portfolio revealed how the effects of downturns extend five or six years into the future. Lastly, simulated downturns were conducted assuming 50 percent and 90 percent levels of coverage with a 10-year downturn duration to illustrate how possible savings targets for the LSRA and other state reserves can be calculated using the results of these analyses.

The results of these analyses are intended to inform the effort to identify a suitable, evidence-based savings target for the state's reserves. However, these analyses are not meant to dictate a target number. In order to make this decision, Wyoming policymakers must address pressing questions regarding the fund's purpose, the level of coverage it should provide, the downturn duration it is meant to cover, and when the state should make withdrawals from the fund.

Technical Appendices

Contents

Appendix A: Data Descriptions

In the attached report, Pew presents its adaptation of the Minnesota methodology for identifying an evidence-based savings target to the state of Wyoming. One key component of this adaptation is the use of state-specific data. In the Minnesota method, analysts with Minnesota Management and Budget utilize national data to approximate the taxable base for their state, to avoid conflating their measures of revenue volatility with volatility that would result from changes in tax policy or rates. Pew opted to draw exclusively on state-specific data sources, as provided by the Wyoming Legislative Service Office (LSO). In many cases, this meant utilizing actual state revenue data with the acknowledgement that measurement error would be introduced due to changes in tax policy; however, it was deemed more appropriate than attempting to estimate the taxable base using national data as Wyoming's revenue and economic experience are so different. Where possible, non-revenue data were utilized to reduce the impact of tax policy changes; however, some of this measurement error is unavoidable given the data available. Unless otherwise noted, the data were provided by the Wyoming LSO.

Sales and Use Taxes

Pew used State of Wyoming data recording actual sales and use tax collections from 1970 through 2014. Note that because real revenues are used, no corrections have been made to account for changes in tax policy, potentially inflating the volatility of the series. However, the impact of this increased volatility is limited, since sales and use tax revenue comprises a relatively small share of Wyoming's overall revenue portfolio.

Efforts were made to estimate state sales and use tax using national Bureau of Economic Analysis (BEA) data adjusted for Wyoming's state-specific sales and use tax exemptions. However, based on discussions with the Wyoming LSO it was decided that the need to modify those data to reflect the impact of mining activity in Wyoming would likely introduce more error than it would correct.

Severance Taxes

Severance tax revenue streams included in this analysis comprise (1) oil, (2) coal, (3) natural gas, and (4) "other" severance tax streams. Note the "other" severance tax stream refers to the taxes generated by the extraction of several less prominent minerals such as uranium and trona. Volatility in this sub-portfolio is evaluated using actual value assessments from 1948 through 2013 and uncertified estimated value assessments for 2014. Value assessments record the full monetary value of the mineral extracted from the state in a given year, while severance tax collections reflect the value assessment multiplied by the tax rate for that given mineral. This approach allows for a full examination of mineral extraction and volatility in possible severance collections, while avoiding any measurement error that would be caused by changes in severance tax rates.

These data were examined as a sub-portfolio, meaning the full Minnesota methodology was performed on the four severance streams in isolation from the other major revenue streams. The data utilized in the full portfolio analysis for "severance taxes" reflect the results of the sub-portfolio analyses and report the conditional standard deviation and standardized residuals of the severance portfolio from 1949 through 2014.

Of note, severance tax assessment changed from being assessed on prior year production to being assessed on current year production in 1981. Prior to 1981 both ad valorem (property) tax and severance tax assessment were based on prior year mineral production. While this reflects an important shift in how the state's severance taxes are collected, Pew's analysis does not incorporate this lag, as any estimated shortfall should be informed by the most current data available. Expectations for any revenue shortfall and state budgeting decisions would draw upon current data, as such, no lag is appropriate.

Investments

Investment income generated for the state is recorded in two streams: (1) pooled general fund investment income, recorded between 1970 and 2014, and (2) earnings from the Permanent Wyoming Mineral Trust Fund between 1976 and 2014. For both streams, the data record actual income generated for the state by its investment holdings. The pooled GF income series records revenue generated by pooled investments and deposited directly into the general fund. The PWMTF series measures investment earnings from the state's mineral trust fund, which is funded by mineral extraction. In both cases, there is no concern for changes in tax policy as these data reflect investment income and are not subject to any changes in tax rate.

Similar to the severance tax streams, these data were examined as a sub-portfolio. Both streams were examined in tandem and combined to form an "investment income" portfolio that is then used as a major tax stream for the full portfolio analysis. The data utilized in the full portfolio analysis for "investment income" reflect the results of the sub-portfolio analyses and report the conditional standard deviation and standardized residuals of the severance portfolio from 1971 through 2014.

Mineral Royalties

The federal mineral royalties revenue streams represent payments made to Wyoming for mineral extraction of federally-owned minerals in the state. These data are recorded in two streams, (1) federal mineral royalties, recorded from 1970 through 2014, and (2) coal lease bonus payments, for the period between 1981 and 2014. Federal mineral royalties record the percentage of total mineral extraction value that is paid to the state by the federal government for mineral extraction of federally-owned minerals within the state. Coal Lease Bonus Payments (CLBP) reflect revenues generated by auctions for access to federally-owned coal resources. Unlike the severance tax streams, data for mineral royalties reflect the raw, actual revenues generated by each stream. This is done for two reasons: 1) value assessments are not available for these streams; and 2) each stream does not follow typical revenue processes, which might inflate the volatility – an aspect of these streams that is important to include in any estimate of the stream's volatility as they are beyond the control of the state.[1]

---

[1] For instance, the Coal Lease Bonus Payments (CLBP) stream reflects revenues generated by auctions for access to federally-owned coal resources. These auctions then provide revenue to the state of Wyoming similar to the more

Similar to the severance and investment streams, these data were examined as a sub-portfolio. Both streams were examined in tandem and combined to form a "mineral royalties" portfolio that is then used as a major tax stream for the full portfolio analysis. The data utilized in the full portfolio analysis for "mineral royalties" reflect the results of the sub-portfolio analyses and report the conditional standard deviation and standardized residuals of the severance portfolio from 1971 through 2014.

K-12 Property Taxes

K-12 property taxes collected reflect the levied collectible in 1990. While this reflects an important shift in how the state's property taxes are collected, Pew's analysis does not incorporate this lag, as any estimated shortfall should be informed by the most current data available. Expectations for any revenue shortfall and state budgeting decisions would draw upon current data, as such, no lag is appropriate. The LSO indicated a high volume of changes in the property tax rates, which influence how the K-12 mill levy is calculated. While these changes technically reflect shifts in tax policy, Pew has deemed these policy changes as being an integral component to the volatility in this stream. Whereas it is appropriate to avoid the inclusion of tax policy changes in most cases, as they reflect an artificial inflation of the variance, the fact these policy changes occur so frequently suggests they are one of the primary contributors to volatility in this stream as opposed to a random event that would bias any estimates of the stream's underlying volatility.

"Other" Revenues

This category reflects state revenues generated by numerous smaller tax and non-tax revenue streams from considerably smaller sources that would not warrant examination or inclusion individually. However, as the combined contribution of these revenue streams nears 10 percent of the full revenue portfolio, it was deemed appropriate for their inclusion as a single major revenue stream. Examples of these smaller revenue sources include but are not limited to: franchise taxes, cigarette taxes, excise tax on wind generated electricity, penalties and interest for unpaid taxes, taxes on services provided, and state royalties.

These data are available from 1980 through 2014, but reflect two separate streams of "other" revenues. From 1980 through 2004, only one stream is recorded; however, starting in 2005, a separate K-12 stream is recorded. Due to the comparatively short duration of the separate K-12 stream, the two "other" revenue streams are combined as a simple sum of revenues collected as opposed to being subject to a sub-portfolio treatment. As such, the "other revenues" stream included as a major stream in the full portfolio analysis reflects the raw collections for the full period, with the caveat that 2005 through 2014 is the sum of both normal collections and K-12 collections.

---

consistent federal mineral royalties stream. However, the CLBP stream is highly erratic, with no set date for the auctions or guarantees regarding the rates, how much access will be leased, or generally, how much revenue they will generate for the state. As such, actual revenues are a more appropriate measure here as they will capture the volatility that results from the unique nature of the revenue stream.

Appendix B: Additional Methodological Details/Descriptions

In this section, Pew presents the specific technical components of its adaptation of the Minnesota methodology for identifying an evidence-based savings target to the state of Wyoming. As documented in the attached report, the methodology proceeds through five key steps: (1) separate the cyclical deviations from observed volatility, (2) identify the time-varying volatility for each revenue stream, (3) calculate the time-varying covariance between revenue streams, (4) estimate the portfolio volatility, and (5) translate the estimate into a savings target. While these steps are explained in the attached report, the more technical components are relegated from the main report to this appendix.

Separate the cyclical deviations from observed volatility

The objective of this step in the methodology is to decompose observed volatility for a given revenue stream into its cyclical and non-cyclical components. The cyclical component, or trend, reflects the volatility in revenues that can be attributed to the normal business or revenue cycle. As a result, this component of the observed volatility is removed from the series to allow for the modeling of all non-cyclical volatility in revenues. This is achieved through a series of steps detailed below.

While nominal dollars are easily the most intuitive measure available, they provide a poor measure for capturing change, as shifts in revenue streams of different scales are not directly comparable. To address this issue, a logarithmic transformation is performed on the raw data series. The justification for this process is rooted heavily in the economics literature where such transformations are both accepted and normal procedure. More specifically, this transformation converts the data in a way that allows for their year-over-year changes to approximate percentage changes while also not encountering issues when balances increase significantly from one year to the next – a problem that is likely to arise in highly volatile revenue streams.[2]

The next step is to once again transform the data, this time to examine year-over-year changes. The intuition behind examining year-over-year changes is relatively straightforward, as we are most interested in measuring volatility in those changes for each revenue stream over time. As stated above, one strength of logarithmic transformations is that they approximate percentage changes when they are first differenced. The resulting data are a set of log-returns recorded over the full period of the revenue series.

Having made these necessary transformations to the raw investment revenue data, it is possible to decompose the cyclical and stochastic components of each revenue stream. As the data are measured

---

[2] For example, if a revenue stream only had $5 million one year and $125 million the next, then that would report a 12,400% increase. For log-transformed data, the same increase would report a more manageable 382.83% increase. Note that the increase is still sizeable, but much more manageable within the context of estimating volatility.

annually, a Hodrick-Prescott (HP) filter is utilized to calculate the cyclical trend for each revenue series.[3] Figures detailing this transformation are presented on page 5 of the primary report or on pages 15 through 26 in Appendix C.

Hodrick-Prescott filters operate using the following formula:

$$\min \sum_{t=1}^{T} (y_t - \tau_t)^2 + \lambda \sum_{t=2}^{T-1} [(\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1})]^2.$$

In this formula, the filter seeks to minimize the squared deviations between the trend, recorded τ_t, and observed values in y_t. The values are squared to account for negative values when the observed data are below trend. The λ parameter can be thought of as a "smoothing" term that applies weight to how the filter weighs deviations from the trend. With a higher λ the trend line becomes more flexible and reactive, while lower λ values result in a smoother trend line. This is evident in applications of the HP filter to different levels of time measurement. For instance, data measured quarterly[4] feature $\lambda = 1600$ while annual level data use $\lambda = 400$.[5] The higher λ in quarterly data allow the filter to be more flexible given the likely higher levels of variation that will be observed across the more finely measured data.[6]

Having calculated the trend components of each series using the HP filter, the final component to Step 1 is to calculate the deviations from trend for each respective series. The justification here is that the stochastic component is what drives revenue volatility and, in turn, is the most reflective measure of volatility for each revenue series. To achieve this, Pew calculates the difference of the observed data point and the estimated trend for each fiscal year. These new difference-from-trend series are the primary data for the modeling in the subsequent steps.

Identify the time-varying volatility for each revenue stream

The next step is to estimate an Integrated Generalized Auto-Regressive Conditional Heteroskedasticity (IGARCH) model on each of the deviations-from-trend series to produce a measure of the time-varying cyclical volatility for each respective series. IGARCH is a relatively unique type of model that operates in a univariate context to produce estimates for the variance of a series, conditioned on the variance observed thus far. In lay terms, this means the model produces a measure of the variance for the series in each fiscal year that is designed to account for the variance we observe in each preceding fiscal years.

---

[3] A lambda value of 400 is used for the HP filter. Alternative filtering options have been explored and it was determined that there was no significant difference across filters or specifications.

[4] This is the most frequently utilized time interval with HP filters.

[5] There are numerous references to varying levels for λ set equal to 10, 40, and 100 for annual level data; however, the 400 value is established both by Stinson and Schoeppner (2014) in the Minnesota methodology and is rooted in the economics literature as an appropriate value for annual data.

[6] To provide additional context, Ravn and Uhlig (2002) state that λ should vary by the fourth power of the frequency observation ratio, so at the monthly level $\lambda = 129{,}600$.

IGARCH models operate by imposing a specific constraint on the more general GARCH model framework that uses the following formula:[7]

$$\sigma_t^2 = a + br_{t-1,t}^2 + c\sigma_{t-1}^2$$

Here the conditional variance, σ_t^2, is a function of a, a constant; r_t^2, the squared mean of the series at time *t-1*; and σ_{t-1}^2, the variance at time *t-1*.[8] The most important point to note is that the standard deviation, our measure of volatility, is influenced by the standard deviation in the previous year. Moreover, this approach allows the current measure of volatility to be conditioned on all previous values while applying the strongest weight to the most recently observed volatility.

These models are applied to each of the revenue series individually to produce a time-varying standard deviation, or measure of volatility, for each series over their full duration. Figures displaying these measures can be seen for the sales and use tax on page 8 of the report or pages 42 through 47 in Appendix C for the remaining revenue streams.

Calculate the time-varying covariance between revenue streams

Having calculated the time-varying conditional standard deviations of each of the revenue streams, the next step is to calculate time-varying correlations to provide a measure of the covariance between each possible combination of streams.[9] A measure of the covariance between revenue streams is required to provide a better understanding of how the series move in relation to one another over time. Moreover, the correlations help illustrate how the investment portfolio's volatility might expand or contract over time.

While covariance is typically produced with a simple formula, in order to generate an estimate for the covariance that also takes into account all prior information is not as straightforward. To do this, Pew applies a two-stage process developed by Engel (2002). The first step generates residuals from the two

---

[7] The integrated component of the model is in reference to the fact that the variance is not assumed to have an underlying mean of zero. Rather, as the coefficients for the conditional mean and conditional variance terms are required to sum to 1, this variation on the model behaves slightly differently. In particular, in a standard GARCH model any "shock" to the variance series decays over time, reverting to a series with a mean equal to zero. For the IGARCH series, the decay parameter is equal to 1 (the sum of the ARCH and GARCH terms), meaning that any "shock" to the series is allowed to persist. Most importantly, this change implies that the series is allowed to trend up or down and is not restricted to having a mean of zero.

[8] In this area, the mean term is typically referred to as the ARCH term and the variance component is referred to as the GARCH term. This is because of the difference between ARCH models, which are used to model volatility but when we have clear expectations about when we would observe higher levels of variance, and GARCH models, which are more agnostic regarding the timing of any changes in the variance or variance clustering.

[9] Note that this results in a significant number of correlations needing to be calculated for each respective revenue stream. Six major revenue streams produce a 6X6 correlation matrix. The diagonal of the matrix reflects each stream's variance and the upper and lower diagonals are symmetric. However, this still requires the full calculation of 15 dynamic correlations.

separate streams' respective IGARCH models and standardizes those into the Z-distribution.[10] The residuals reflect how far the model predictions differ from the observed values. The residuals for each series are then standardized by dividing them by their respective standard deviations. The reason behind this is to provide measures of the error for each series that are standardized in a way to allow comparisons across revenue streams. This is key, as the measures of time-varying covariance are recorded as the dynamic correlations between the standardized residuals.

This approach capitalizes on an interesting mathematical property of models and their residuals. Namely, the correlation between two data series matches the correlation between their standardized residuals after having been estimated using the same model. This allows us to make the comparison between the volatilities of each respective series using standardized values, as correlations are sensitive to the magnitude of movement in their respective variables, meaning differences in scaling could result in biased measures of the correlation.

The second step calculates the covariance of the standardized residual series and divides the covariance by the product of the two series' respective conditional standard deviations to produce a time-variant conditional correlation.[11] This is done using the following formula:

$$\rho_{ij,t} = \frac{q_{ij,t}}{q_{ii,t}q_{jj,t}}$$

where $\rho_{ij,t}$ is the correlation between revenue streams *i* and *j* at time *t*, and $q_{ii,t}$ and $q_{jj,t}$ are the standard deviations of revenue streams *i* and *j* at time *t*, respectively. Figures reflecting these correlations between Wyoming's sales and use tax and the other major revenue streams are visible on page 9 of the report or pages 48 through 53 in Appendix C for the other major revenue streams.

Estimate the portfolio volatility

The next step is to produce an estimate for the full portfolio. This is done using a modified version of the Markowitz portfolio choice model. This model uses the following formula:

$$\sigma_{p,t} = \sqrt{\sum_{i=1}^{N}\sum_{j=1}^{N} w_{i,t}w_{j,t}\rho_{ij,t}\sigma_{i,t}\sigma_{j,t}}$$

[10] Note the Z-distribution is normally distributed with a mean zero and standard deviation of 1. $Z \sim N(0,1)$.

[11] The specific dynamic correlation coefficient models are renowned for their difficulty to estimate and the failures of their log-likelihoods to converge. In instances where the model achieved convergence Engel's DCC is utilized. However, for the pairings where the model did not converge, a rolling dynamic correlation was calculated. Upon close comparison, the two measures are very similar, especially as additional observations are added to each series.

where $w_{i,t}$ and $w_{j,t}$ are the weights for each respective tax stream at time *t;* $\rho_{ij,t}$ is the conditional correlation between tax streams *i* and *j* at time *t;* and $\sigma_{i,t}$ and $\sigma_{j,t}$ are the conditional standard deviations for each respective tax stream at time *t*.

Having generated measures of the time-varying volatility for each stream and a measure of the time-varying correlation between streams, the only remaining task is to calculate the weights to assign to each tax stream. To calculate the weights, Pew examined the actual revenue collections for all streams as they relate to the full revenue portfolio. In any given year, the weight for each tax stream is the proportion of the portfolio revenues that can be attributed to that stream.[12] For instance, in 2014, the portfolio revenue is $4.6 billion, and severance taxes account for $960 million, or 21.0 percent of all revenues. As such, the weight for $w_{sev,2014} = 0.21$ for the calculation of the portfolio in 2014. From this point, the construction of the revenue portfolio is a matter of substituting the appropriate values for each respective tax stream, stream pairing, and year and then completing the double summation. The full revenue portfolio volatility is presented graphically on page 8 of the report and is also illustrated in more detail on pages 41 and 42 of Appendix C.

It is worth noting that this portfolio modeling is applied at two different stages in Pew's adaptation of the Minnesota methodology. First, individual revenue streams are used to compile sub-portfolios for the investment income, severance taxes, and mineral royalties major tax streams. Second, the characteristics of the major tax streams are combined to produce the full revenue portfolio as described above.

Translate the estimate into a savings target

Having calculated the full portfolio volatility of 19.1 percent of portfolio revenues in 2014, the final component of Pew's analysis is to convert the volatility estimate into a savings target for Wyoming's LSRA. In the original Minnesota methodology, this is done by first multiplying the volatility estimate by a critical value for the level of coverage determined by the legislature and then multiplying that estimate by $\sqrt{2}$ to account for the biennium budgetary structure. While this conversion methodology works for Minnesota, Wyoming still has yet to determine their desired level of protection or their desired shortfall duration. However, components of the Minnesota approach can still be adapted to these analyses.

The first step of this conversion requires the use of the portfolio standard deviation to create a probability density function for possible estimated shortfalls. This is done by assuming shortfalls to follow a standard normal distribution. Having made this assumption, it is possible to identify the critical value multiplier for the desired level of protection by referencing the Z-distribution critical values for a left-side one-tailed test, which assesses the probability that an observed value falls to the left of the estimate. However, it is critical to recognize the distribution being estimated reflects all deviations from trend including both negative values (downturns) and positive values (surpluses). Since a savings target

---

[12] This requires examination of actual revenues collected, not value assessments or other proxied values.

is determined by the number of downturns the state wishes to offset with their reserves, a slight adaptation of critical values is required, where the desired level of coverage must be determined relative to the half of the distribution attributed to downturns. For example, the critical value for a one-tailed 90% test is 1.28; however, 50 percent of the distribution being examined reflects surpluses. Therefore, it is necessary to adjust the discussion to reflect that the 90% critical value is appropriate for 10% of the full distribution, but that share represents 20% of all possible downturns. As a result, the critical values for a one-tailed 90% test appropriate for providing coverage for 4 out of every 5 possible downturns. Alternately, the higher critical value of 1.65 for a one-tailed 95% test could be used to achieve a 90% level of coverage or enough to protect against 9 out of 10 possible revenue downturns. As Wyoming has not designated a desired level of coverage, Pew identified several critical values for possible levels of coverage at 50, 80, and 90 percent.[13]

In addition to the level of coverage, the state must identify its desired shortfall duration. The Joint Revenue Committee expressed an interest in a 10-year protection window which was used for these analyses. The conversion of the shortfall estimate to cover multiple years was achieved through a thorough examination of the portfolio revenues from 1970 through 2014. Specifically, an auto-distributed lag (ADL) model was estimated on actual revenue collections to estimate the persistence of any revenue shortfalls or surpluses in the observed data. Simply put, the model asks the question: When the state's revenue portfolio experiences a shock, how long does it take before revenues return to normal?[14] The ADL is estimated using the following formula:

$$\widehat{Rev}_t = \alpha + \beta_1 Rev_{t-1} + \beta_2 Trend_t + \varepsilon_t$$

where $\widehat{Rev}_t$ is the estimated revenues at time *t*; α is a constant determined by the model; Rev_{t-1} is the observed revenues from the previous year (time *t-1*); $Trend_t$ is the revenue cycle trend at time *t*; and ε_t is a stochastic error term. The most important component of this model is represented in β_1 which can be thought of as the decay parameter or, alternately, as the persistence of prior observed values in current revenues. Graphical representations of this decay can be seen on pages 59 and 60 of Appendix C.

The model results in an estimated decay/persistence parameter of 0.56, meaning that 56 percent of any revenue shortfall would remain in the following year. For instance, if there is a $100 million revenue shortfall relative to trend in 2014, the model estimates that $56 million will still remain in 2015, suggesting that to achieve full coverage against the shortfall for 2014 and 2015 would require $156 million in savings.[15] Further, this decay parameter operates geometrically, meaning that 56 percent of the shortfall persists in the next year, while only 31 percent remains two years out, and 18 percent three

[13] The critical values in question here are as follows: 50% = 0.68; 80% = 1.28; and 90% = 1.65.
[14] The term "normal" refers to the statistical mean in this context.
[15] Please note this is only an example meant to illustrate how the persistence parameter from the model operates and can be interpreted. The example is not meant to reflect any actual revenue shortfall.

years out. The estimated shortfall reaches near zero, as indicated by less than 5 percent of the initial shortfall remaining, 6 years after the initial event.

Using the decay/persistence parameter, Pew conducted several simulations for Wyoming's revenues given the estimated portfolio volatility of 19.1 percent and converting the estimated percentages of portfolio revenues using the 2014 portfolio total of $4.6 billion. These simulations are presented on pages 12 and 13 of the report and page 61 and 62 of Appendix C. Pew stresses these are only simulations for how much savings would be needed under the assumed conditions and that Wyoming should work to identify the level of coverage and shortfall duration that is appropriate for their needs.

Appendix C: Supplemental Figures for Revenue Streams

The figures for Wyoming's sales and use taxes are presented in the report. However, identical analyses have been completed for each respective revenue stream. These figures and more are presented in the following pages.

Revenue Collection and Data Input Descriptive Figures

Figure 1 – Wyoming's Revenue Portfolio: Shares of Total Revenues, 1970-2014

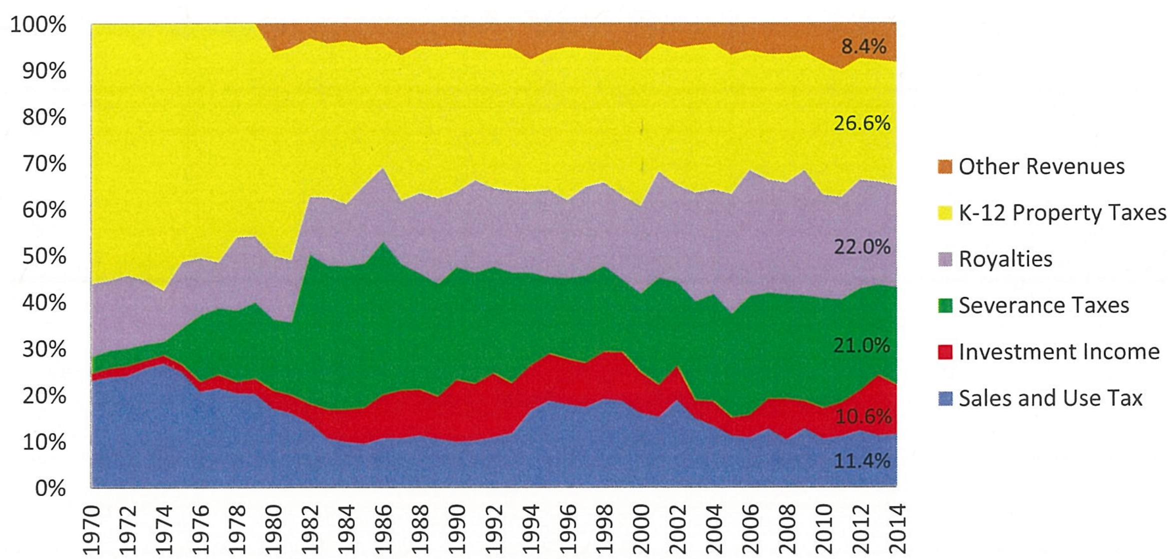

Figure 2 – Wyoming's Portfolio Revenues Collected, 1970-2014

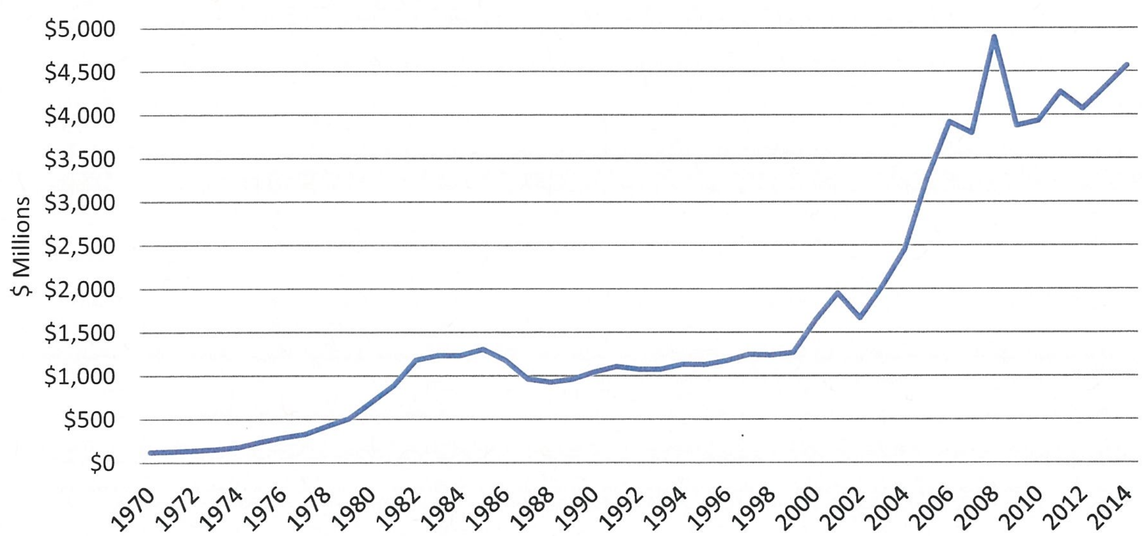

Figure 3 –Sales and Use Taxes Collected, 1970-2014

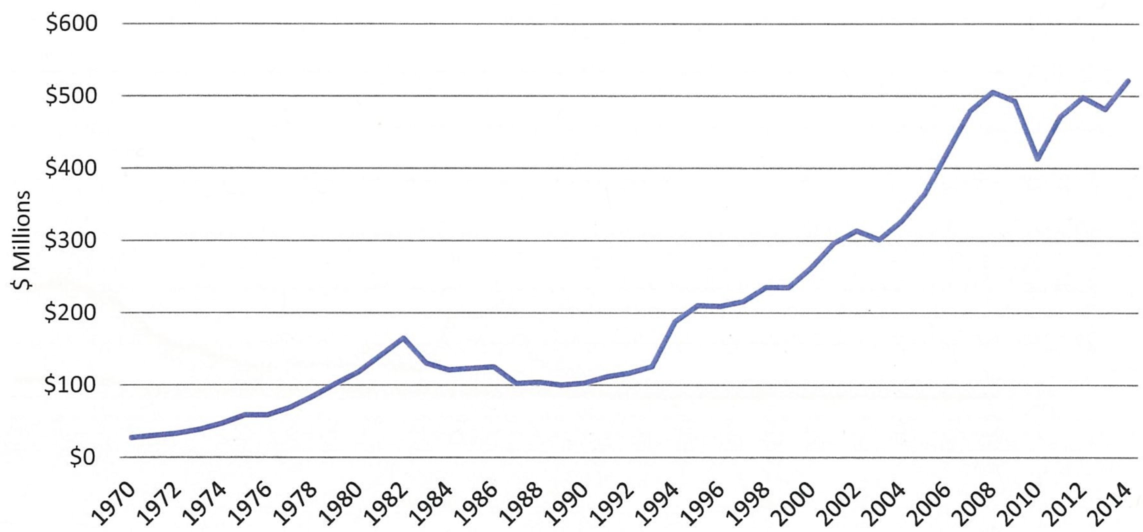

Figure 4 – Investment Income, 1970-2014

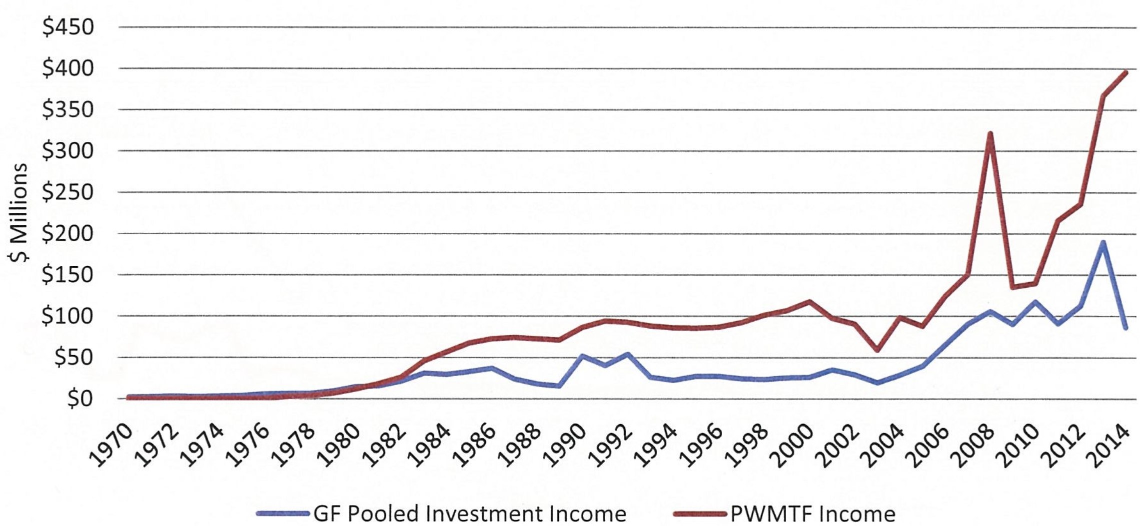

Figure 5 – Severance Value Assessments, 1948-2014

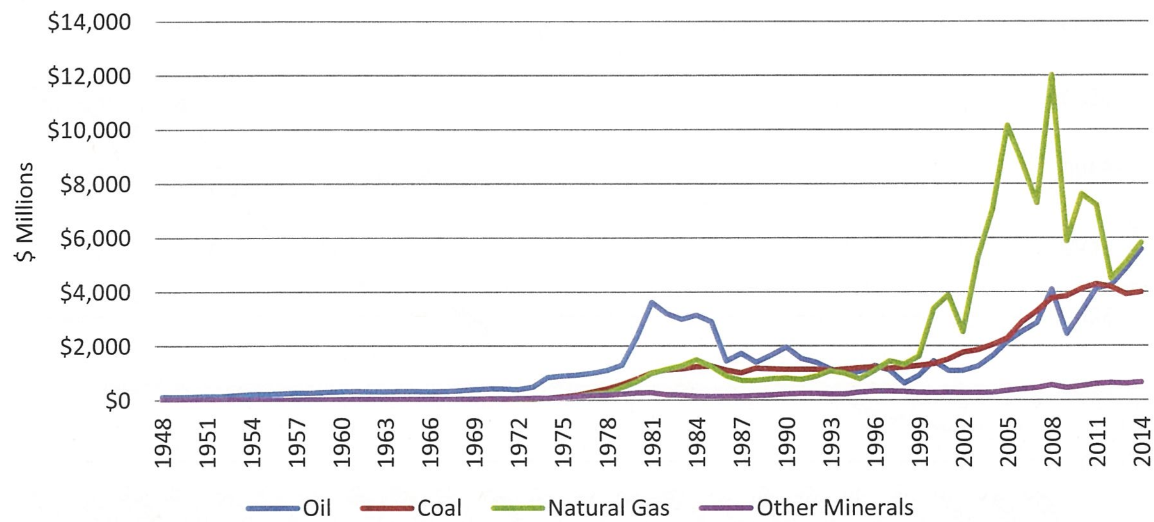

Figure 6 – Mineral Royalties Income, 1970-2014

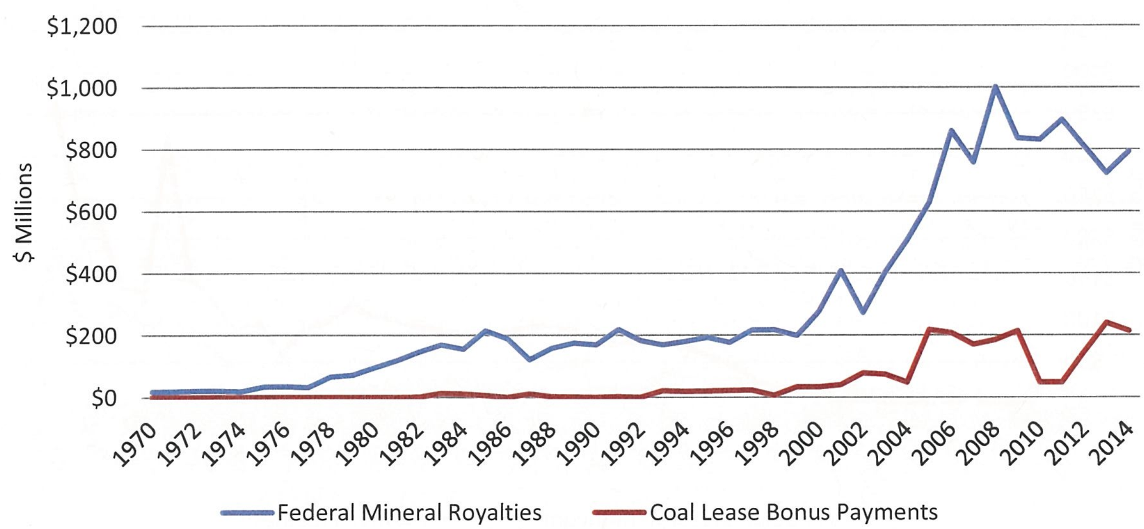

Figure 7 – K-12 Property Taxes Levied, 1948-2014

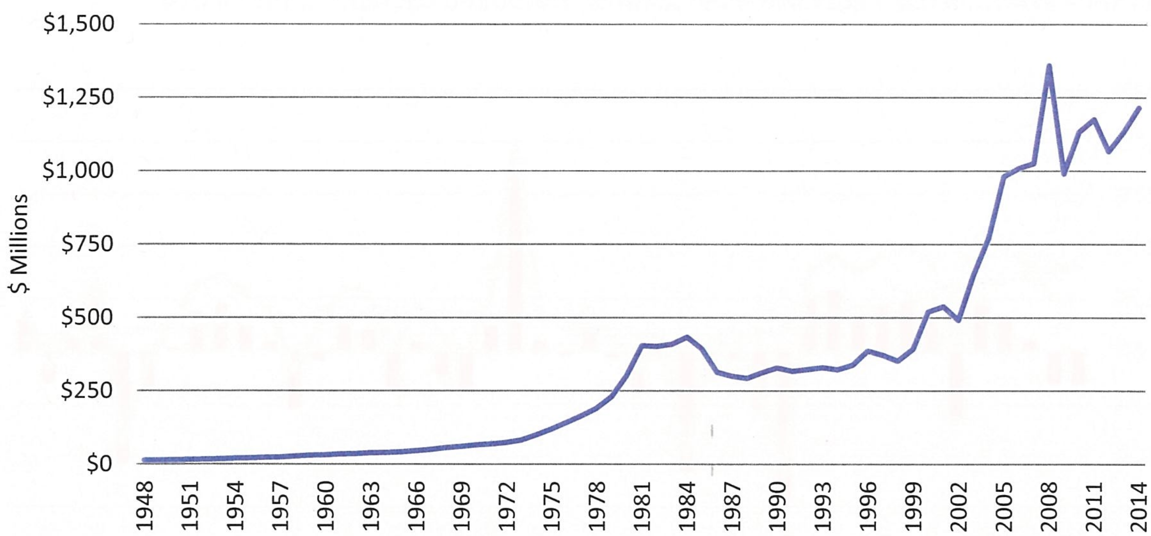

Figure 8 – Revenue from “Other” Sources, 1980-2014

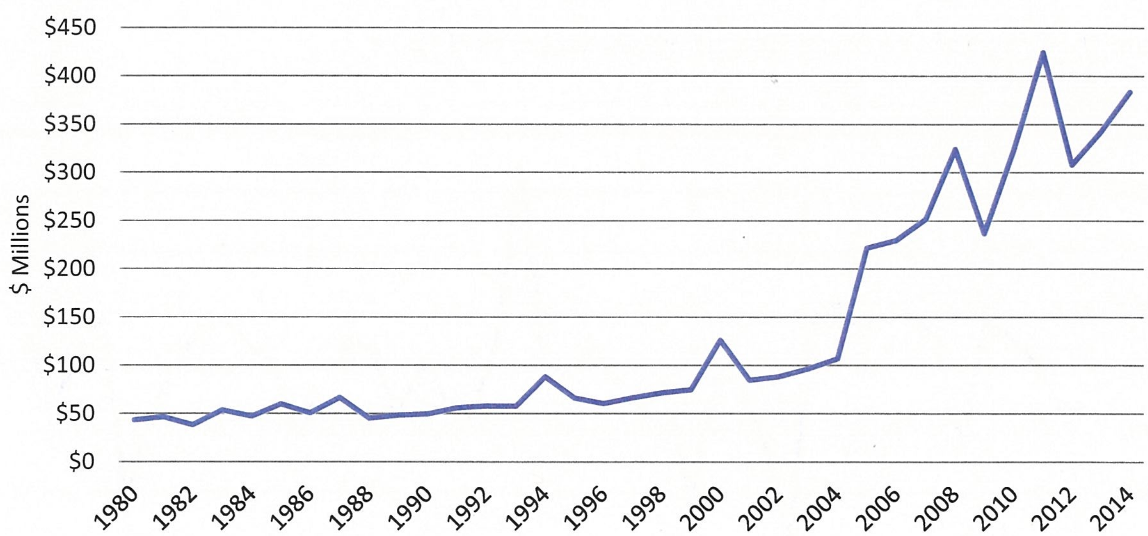

Changes in Revenues, Trend Calculations, and Deviations from Trend

Figure 9 – Sales and Use Taxes: Observed Returns, Trend, and Deviations, 1971-2014

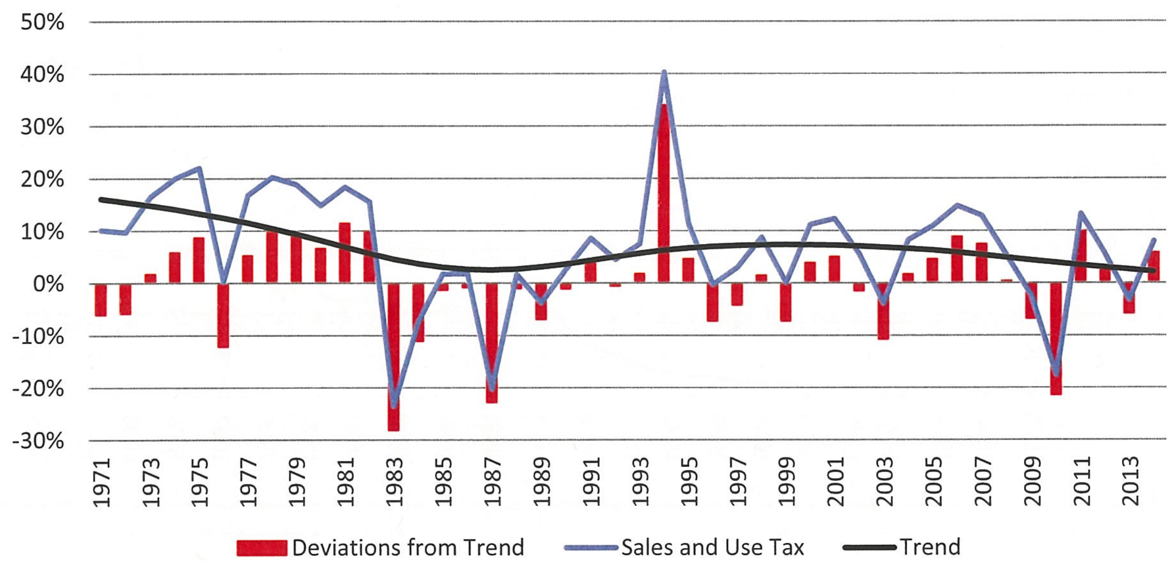

Figure 10 – Sales and Use Taxes: Deviations from Trend, 1971-2014

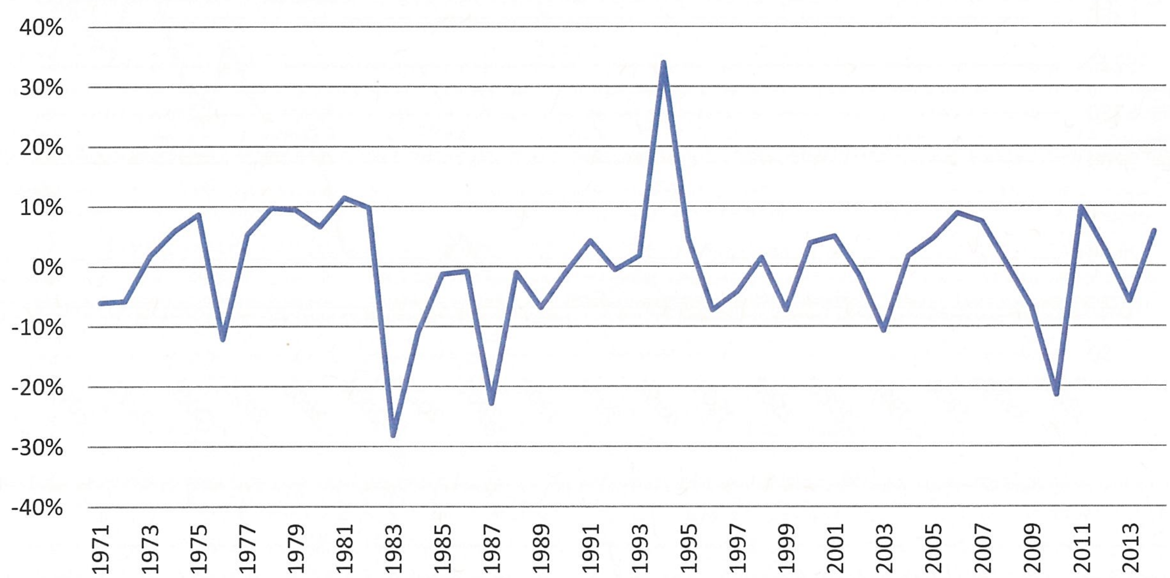

Figure 11 – GF Pooled Investments: Observed Returns, Trend, and Deviations, 1971-2014

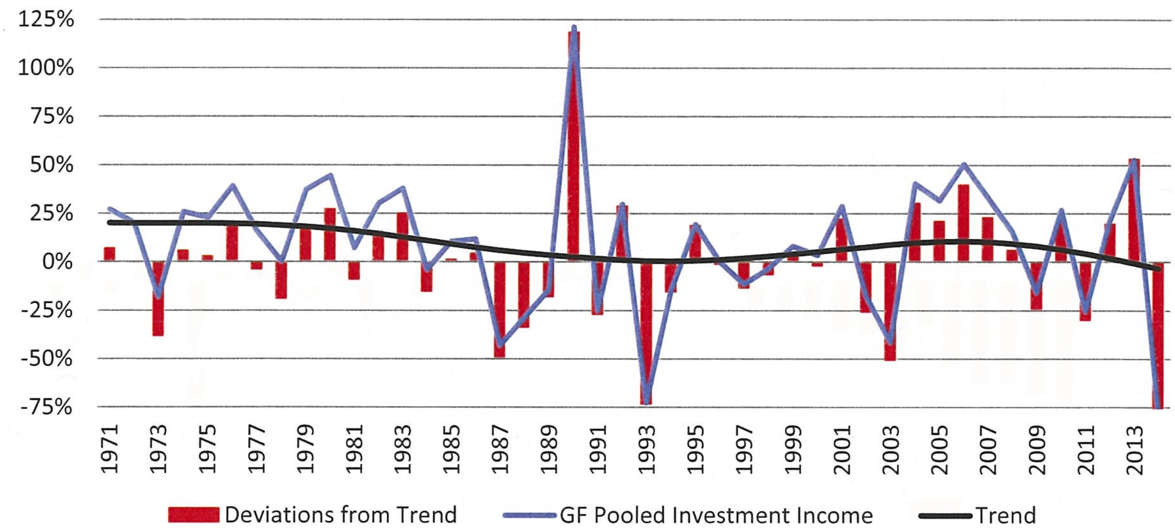

Figure 12 – GF Pooled Investments: Deviations from Trend, 1971-2014

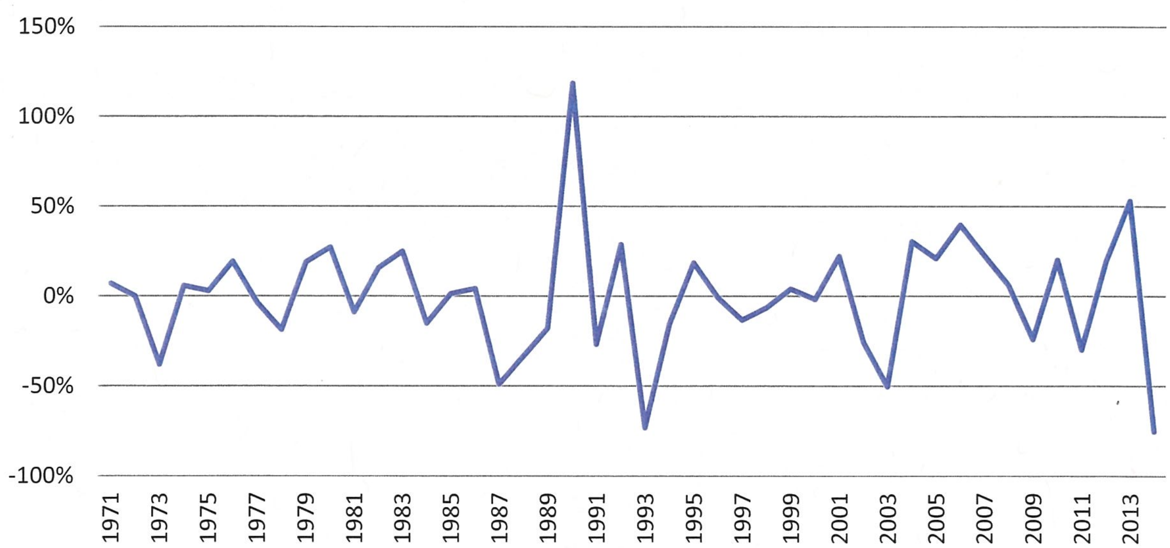

Figure 13 – PWMTF Investments: Observed Returns, Trend, and Deviations, 1977-2014

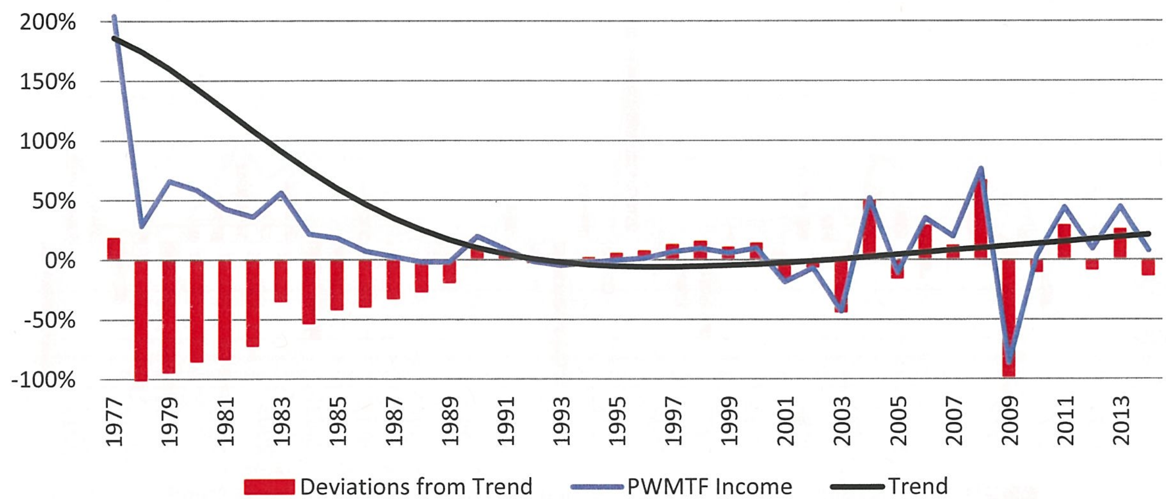

Figure 14 – PWMTF Investments: Deviations from Trend, 1977-2014

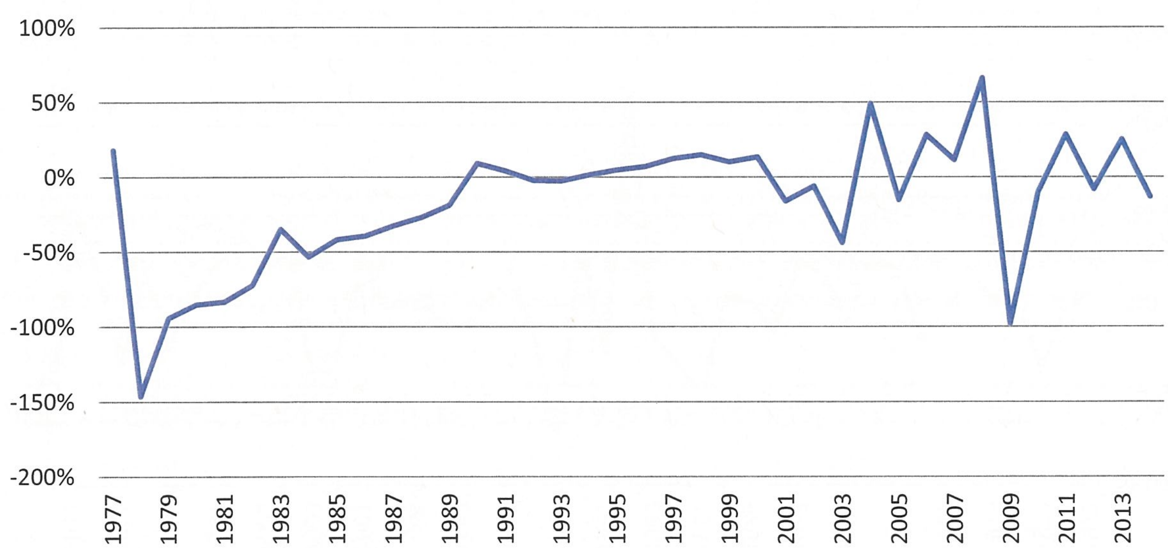

Figure 15 – Oil Assessments: Observed Returns, Trend, and Deviations, 1949-2014

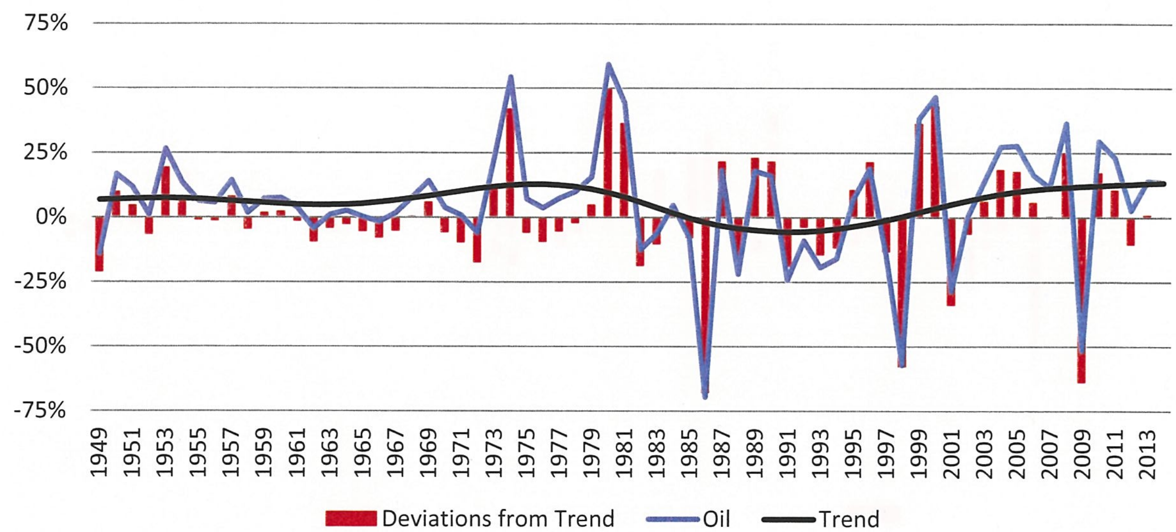

Figure 16 – Oil Assessments: Deviations from Trend, 1949-2014

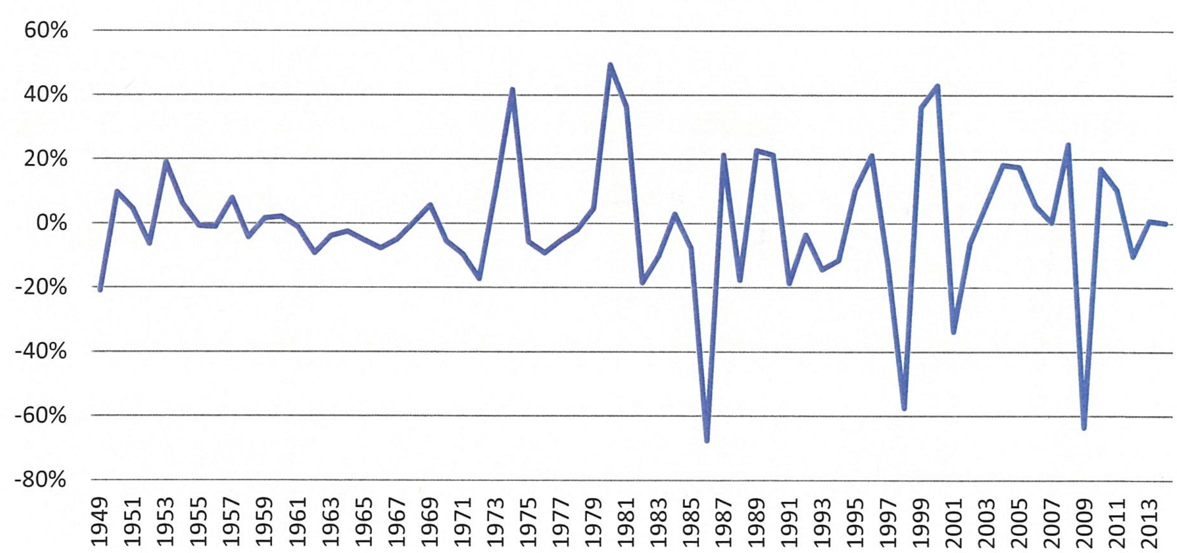

Figure 17 – Coal Assessments: Observed Returns, Trend, and Deviations, 1949-2014

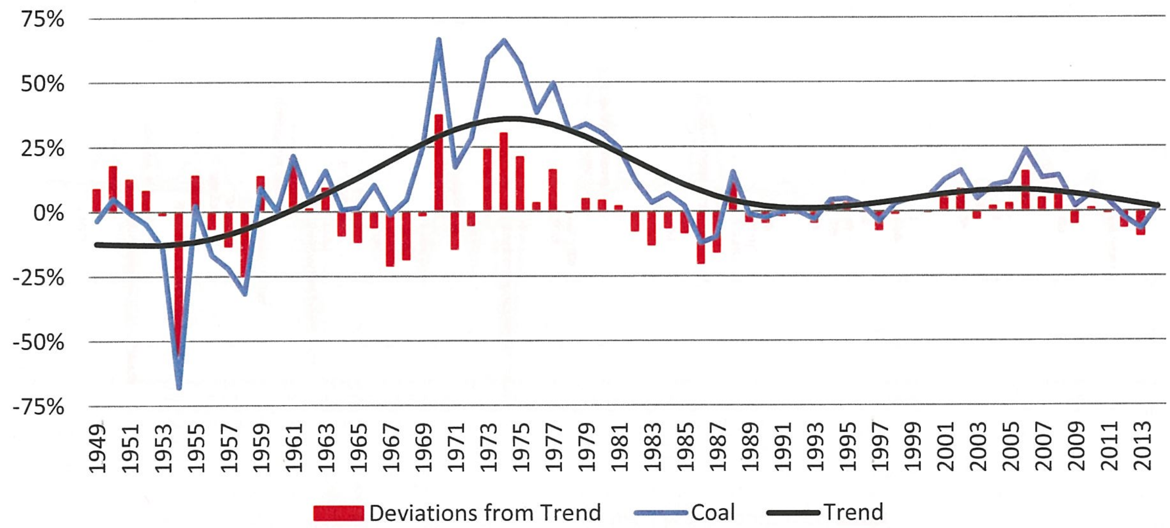

Figure 18 – Coal Assessments: Deviations from Trend, 1949-2014

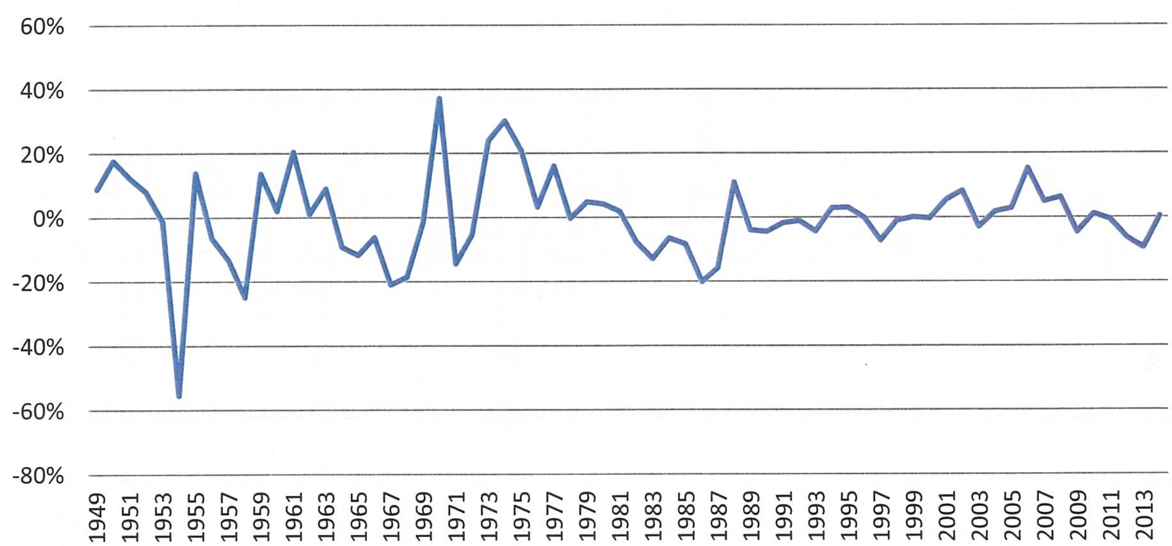

Figure 19 – Natural Gas Assessments: Observed Returns, Trend, and Deviations, 1949-2014

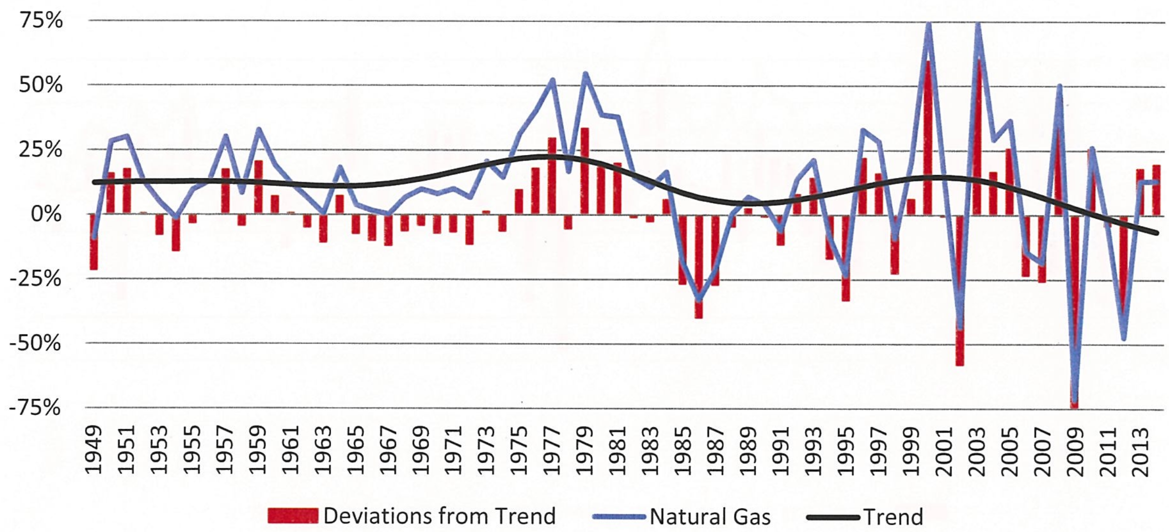

Figure 20 – Natural Gas Assessments: Deviations from Trend, 1949-2014

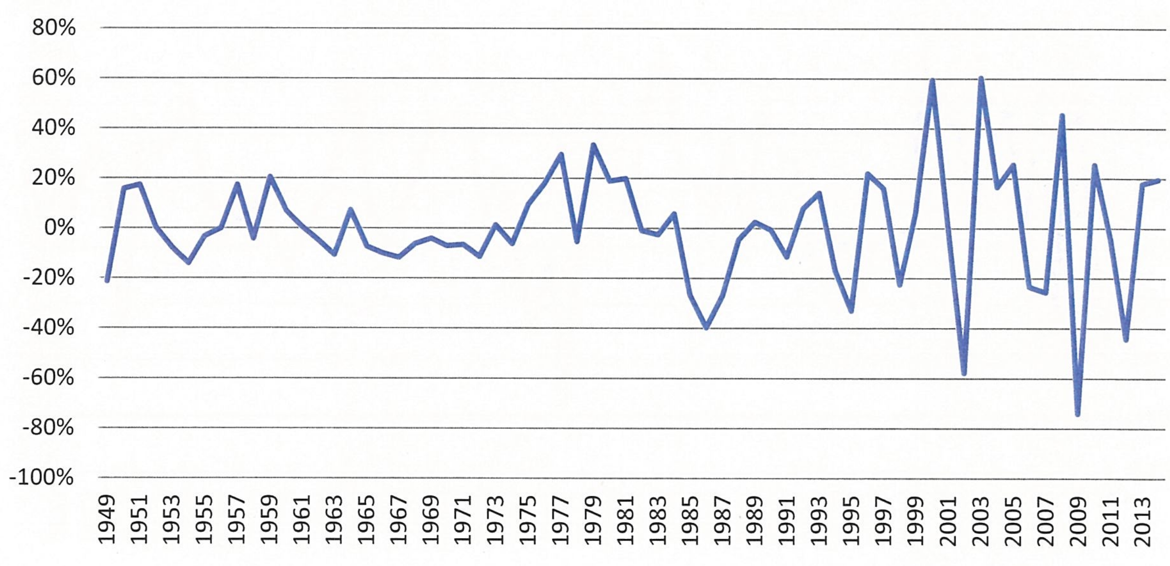

Figure 21 – Other Assessments: Observed Returns, Trend, and Deviations, 1949-2014

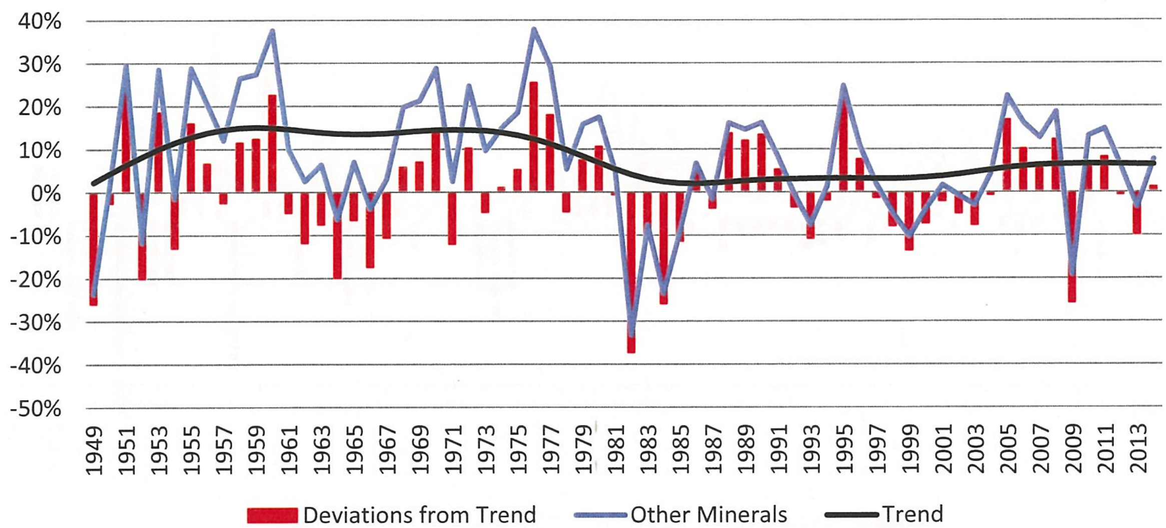

Figure 22 – Other Assessments: Deviations from Trend, 1949-2014

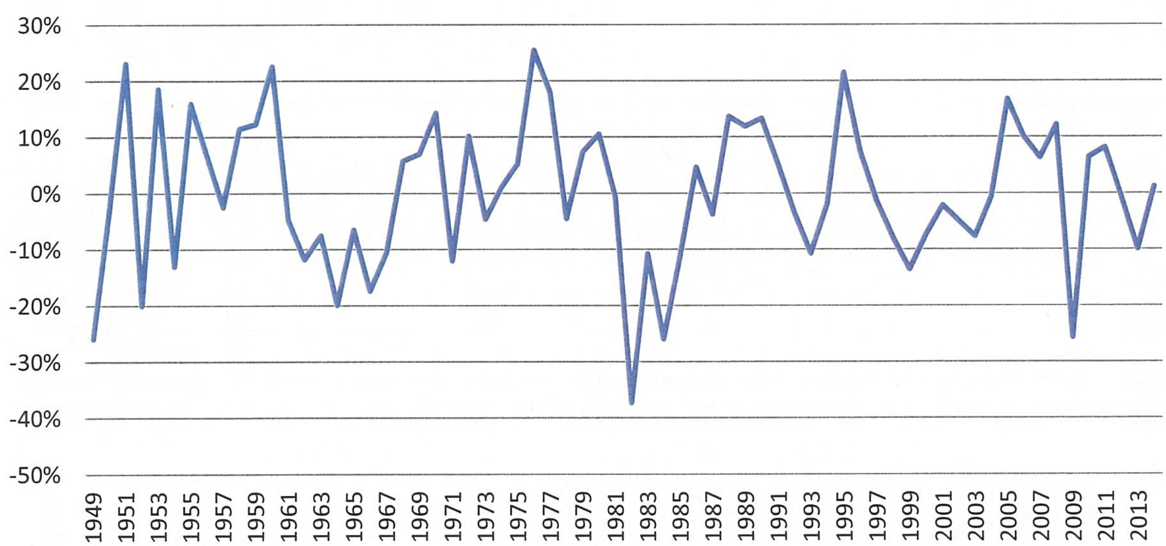

Figure 23 – Federal Mineral Royalties: Observed Returns, Trend, and Deviations, 1971-2014

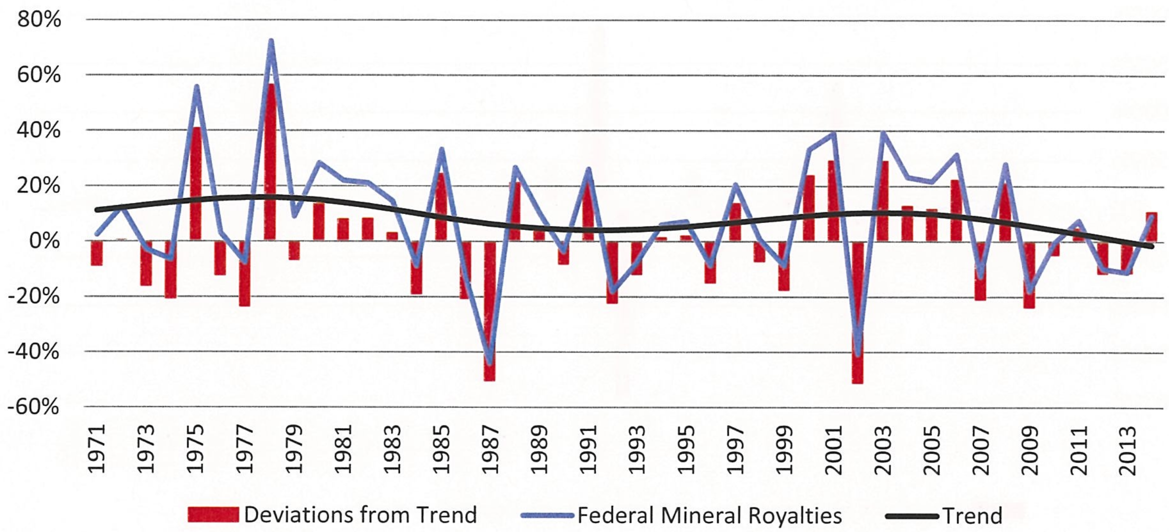

Figure 24 – Federal Mineral Royalties: Deviations from Trend, 1971-2014

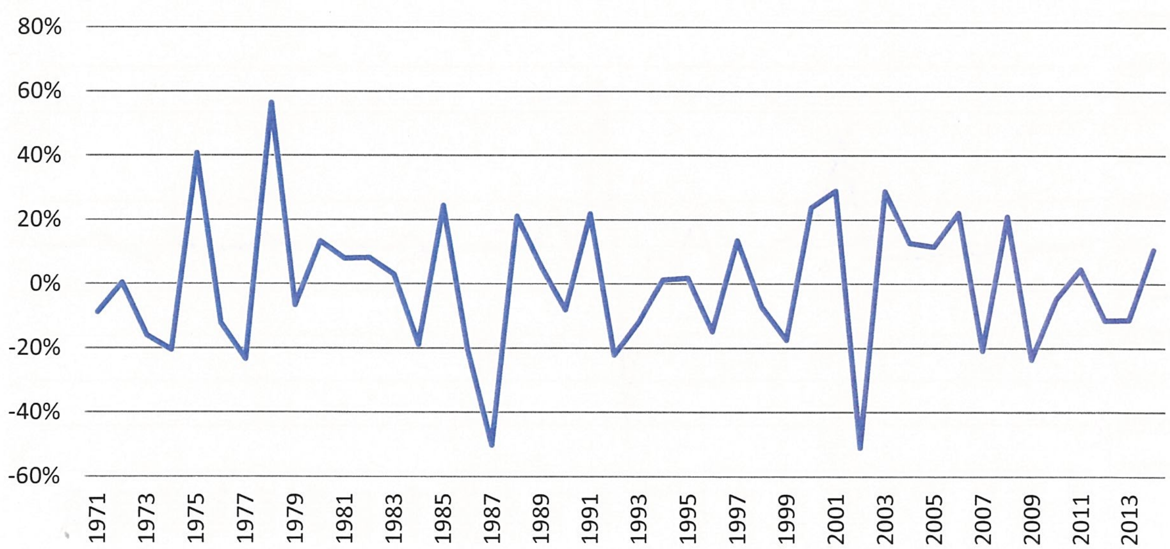

Figure 25 – Coal Lease Bonus Payments: Observed Returns, Trend, and Deviations, 1971-2014

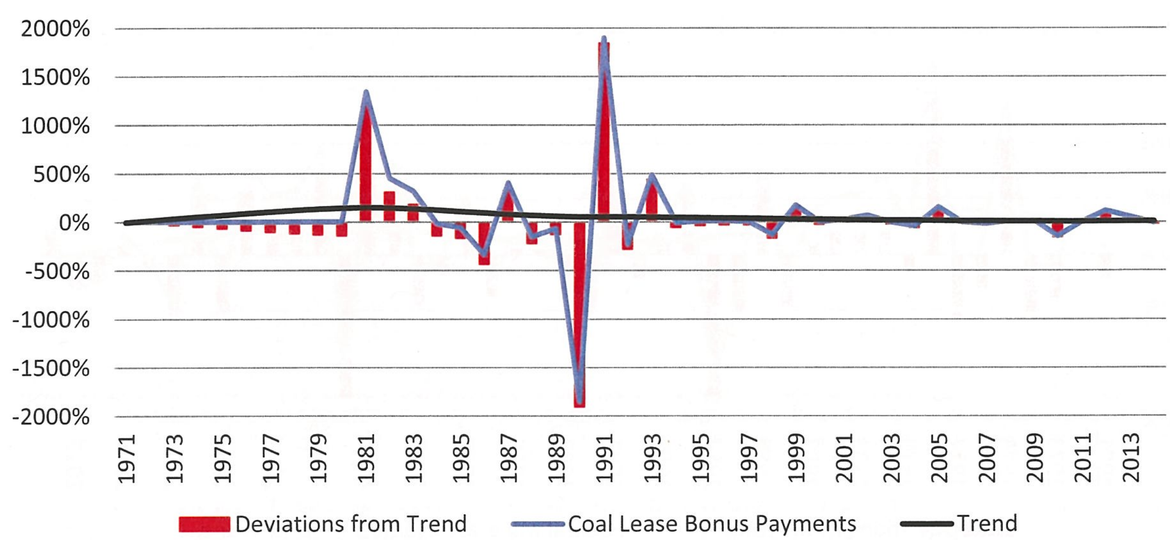

Figure 26 – Coal Lease Bonus Payments: Deviations from Trend, 1971-2014

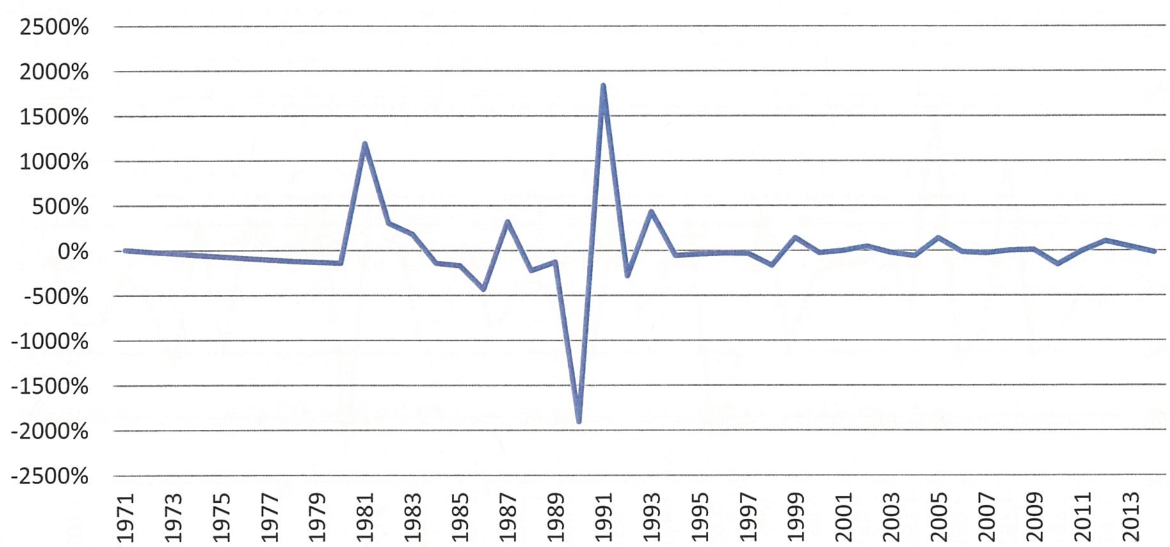

Figure 27 – K-12 Property Taxes: Observed Returns, Trend, and Deviations, 1949-2014

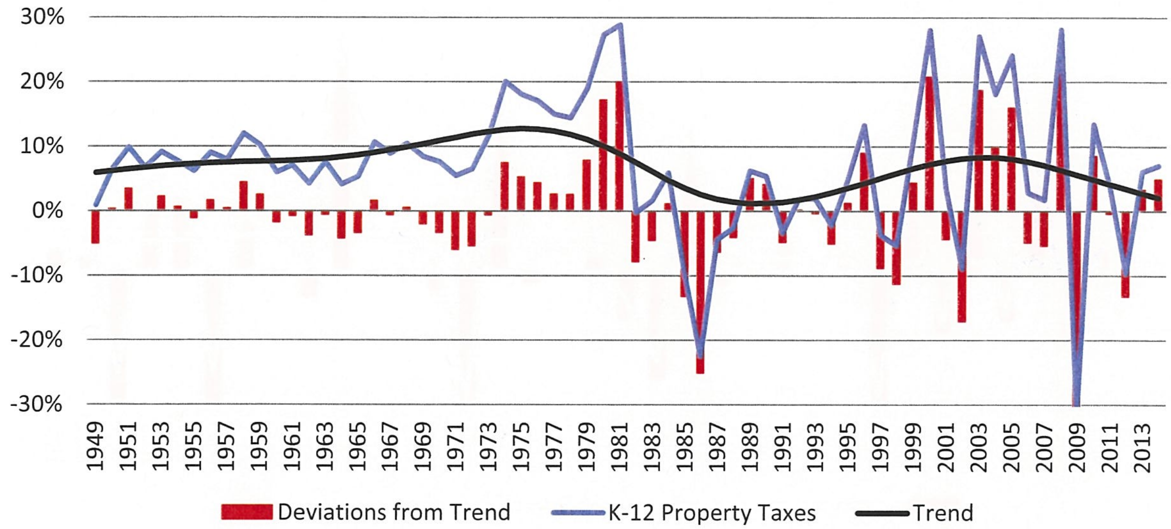

Figure 28 – K-12 Property Taxes: Deviations from Trend, 1949-2014

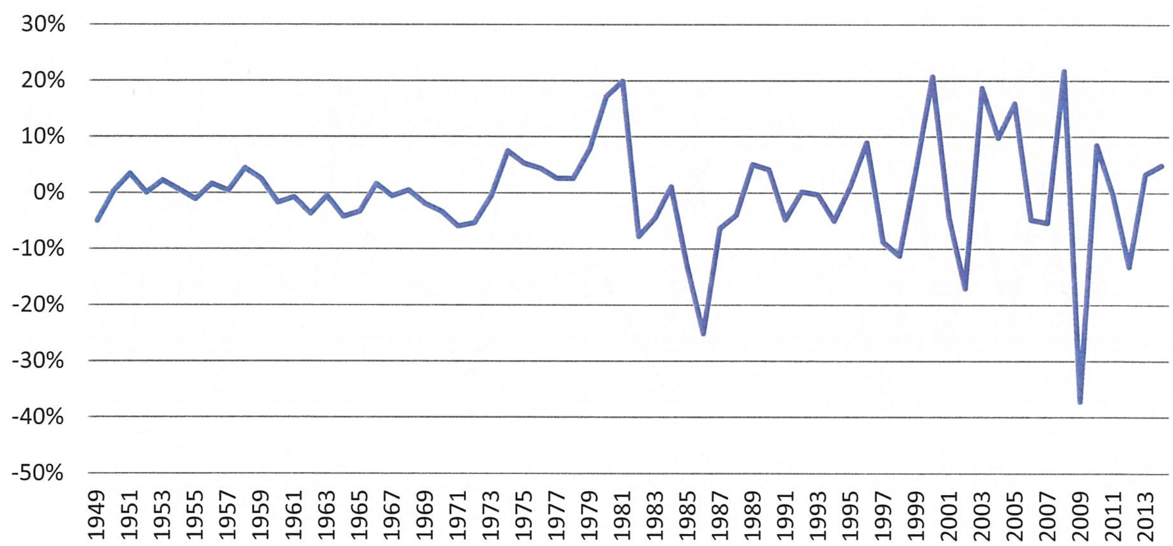

Figure 29 – "Other" Revenues: Observed Returns, Trend, and Deviations, 1981-2014

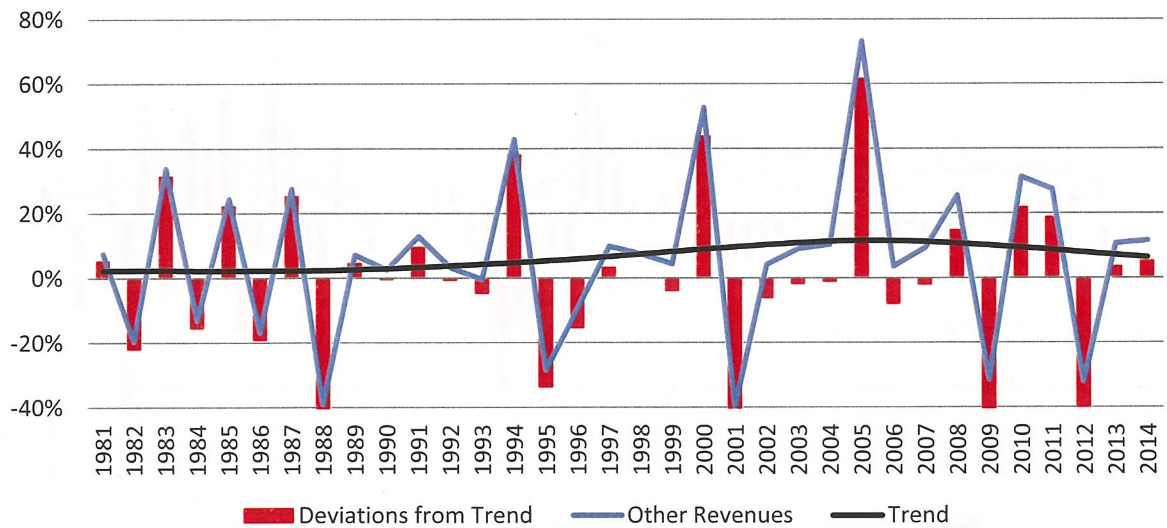

Figure 30 – "Other" Revenues: Deviations from Trend, 1981-2014

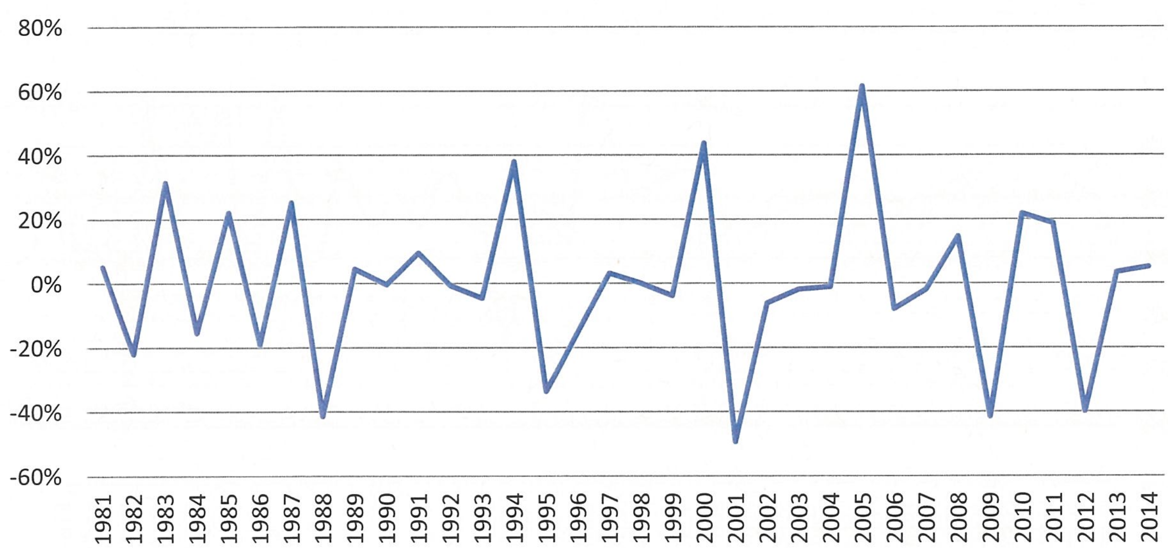

Time-Varying Standard Deviations: A Measure of Volatility

Figure 31 – Characteristics of Sales and Use Taxes, 1971-2014

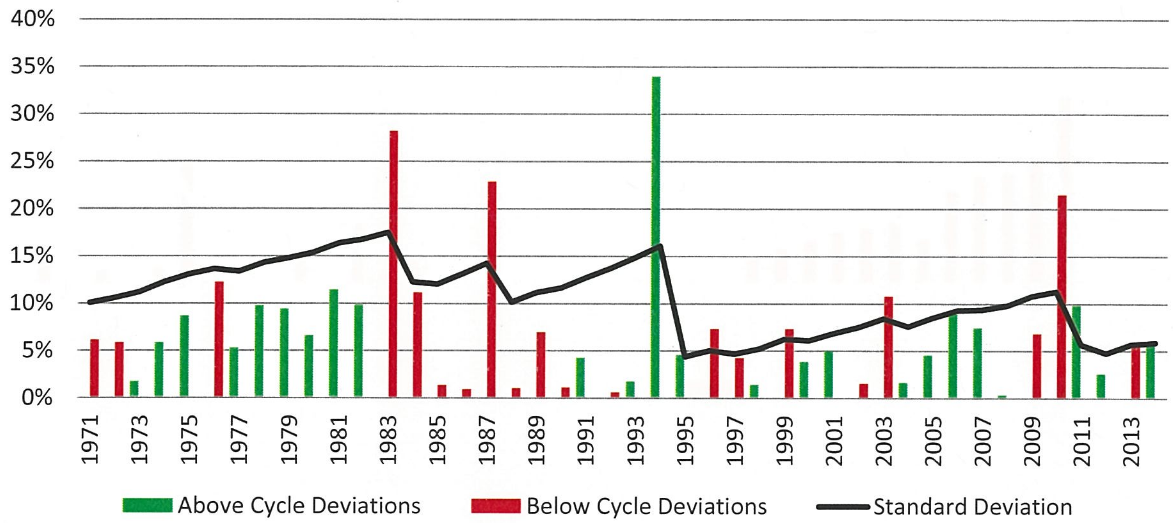

Figure 32 – Characteristics of GF Pooled Investment Income, 1971-2014

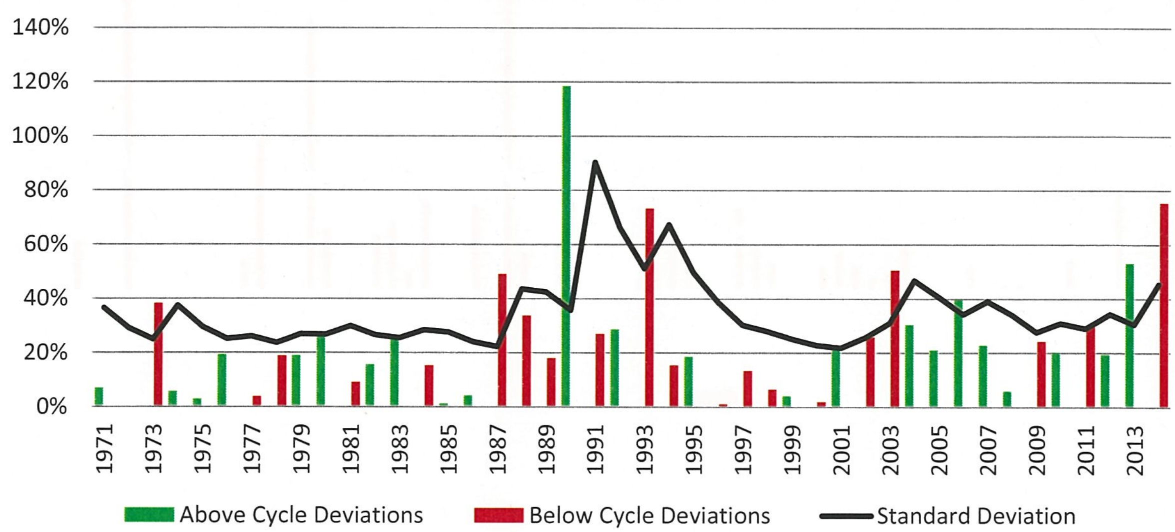

Figure 33 – Characteristics of PWMTF Investment Income, 1977-2014

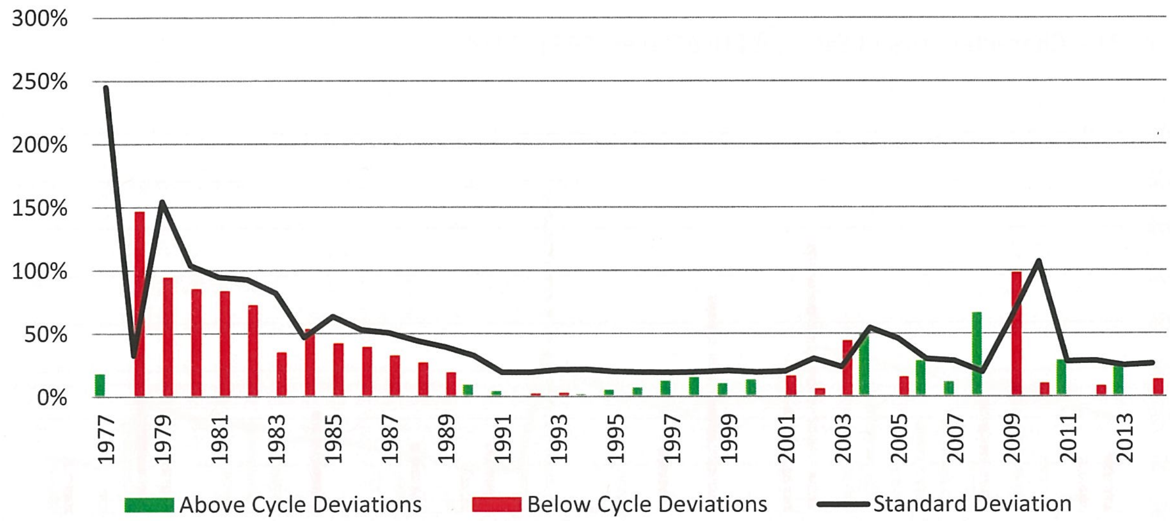

Figure 34 – Characteristics of Oil Assessments, 1949-2014

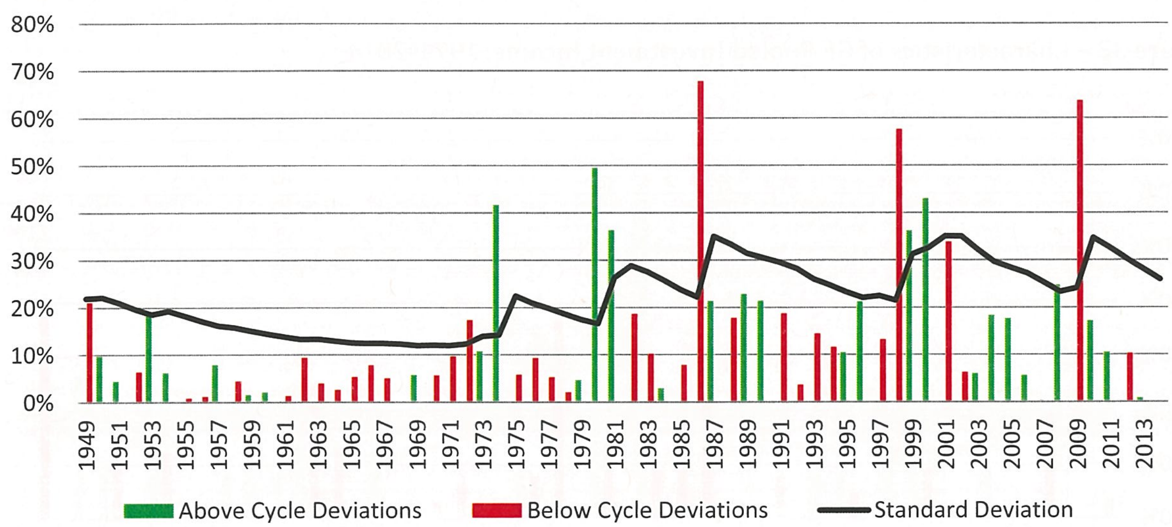

Figure 35 – Characteristics of Coal Assessments, 1949-2014

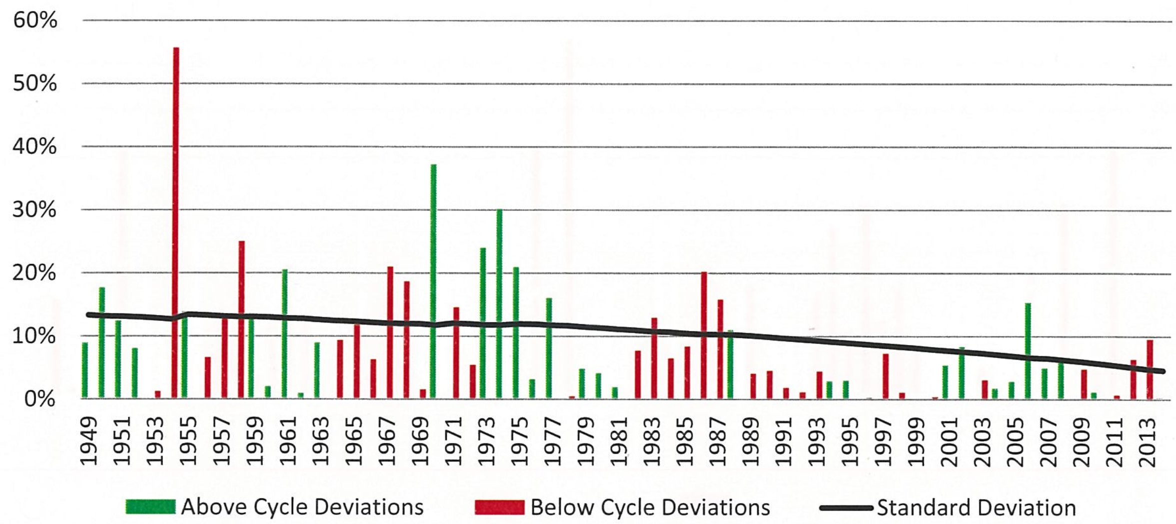

Figure 36 – Characteristics of Natural Gas Assessments, 1949-2014

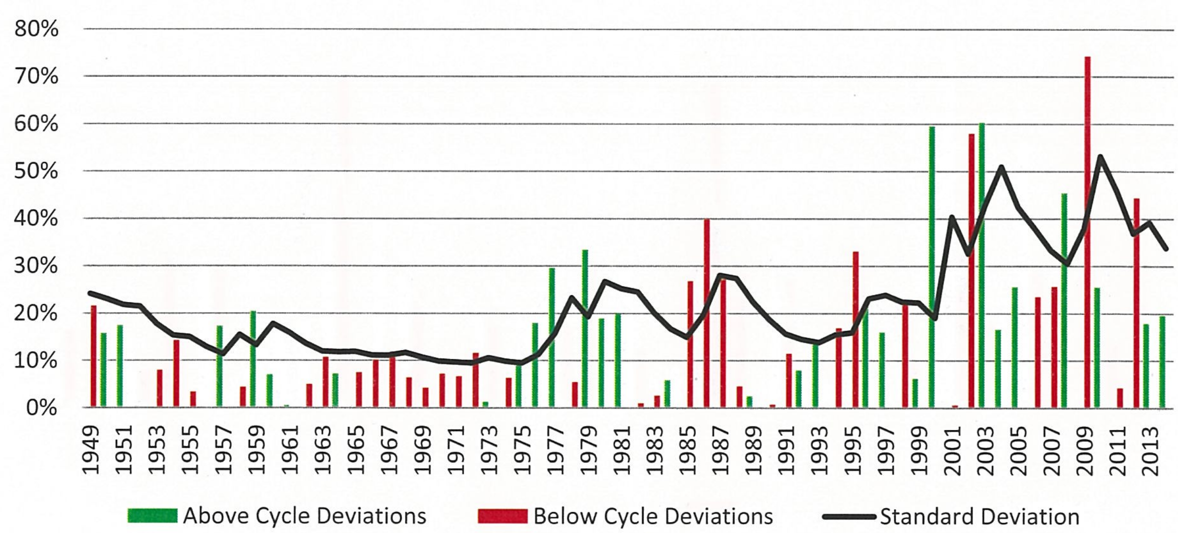

Figure 37 – Characteristics of Other Mineral Assessments, 1949-2014

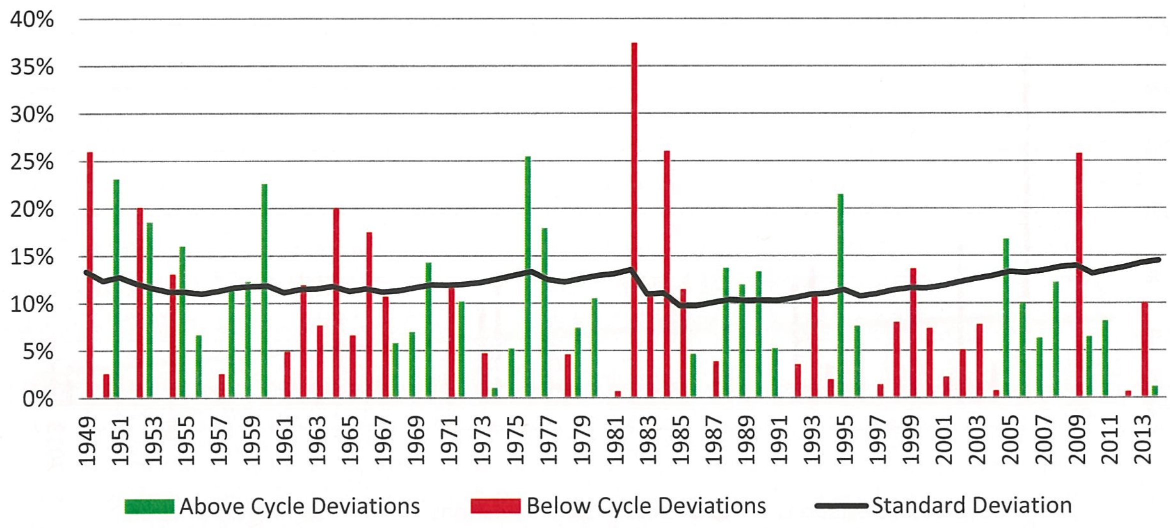

Figure 38 – Characteristics of Federal Mineral Royalties, 1971-2014

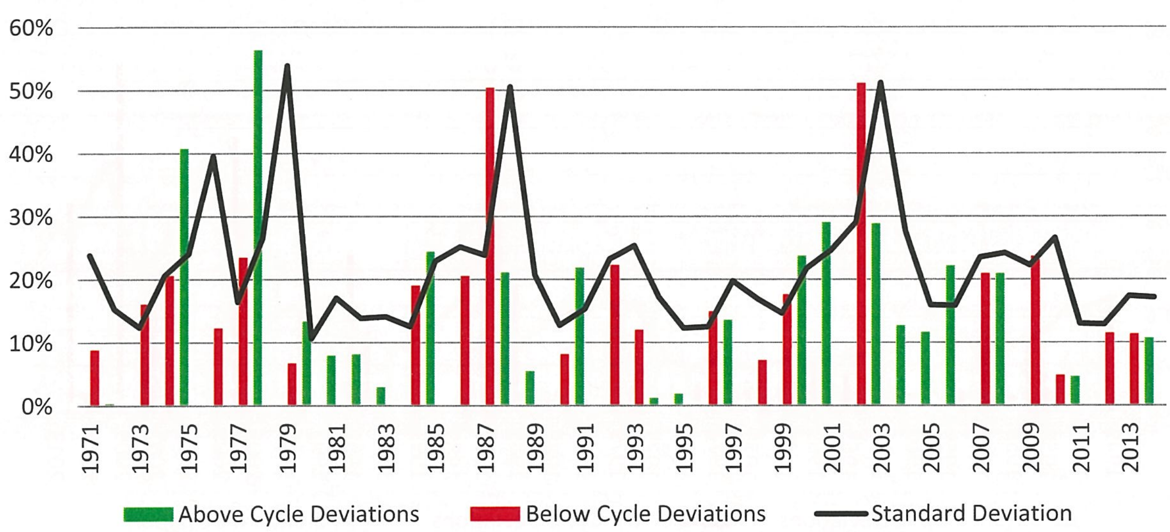

Figure 39 – Characteristics of Coal Lease Bonus Payments, 1981-2014

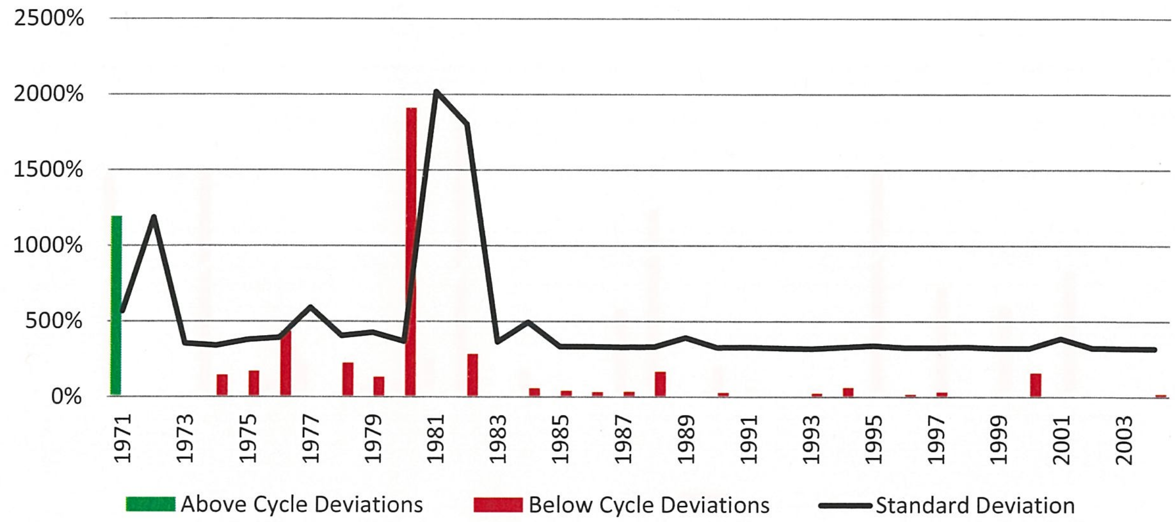

Figure 40 – Characteristics of K-12 Property Taxes, 1949-2014

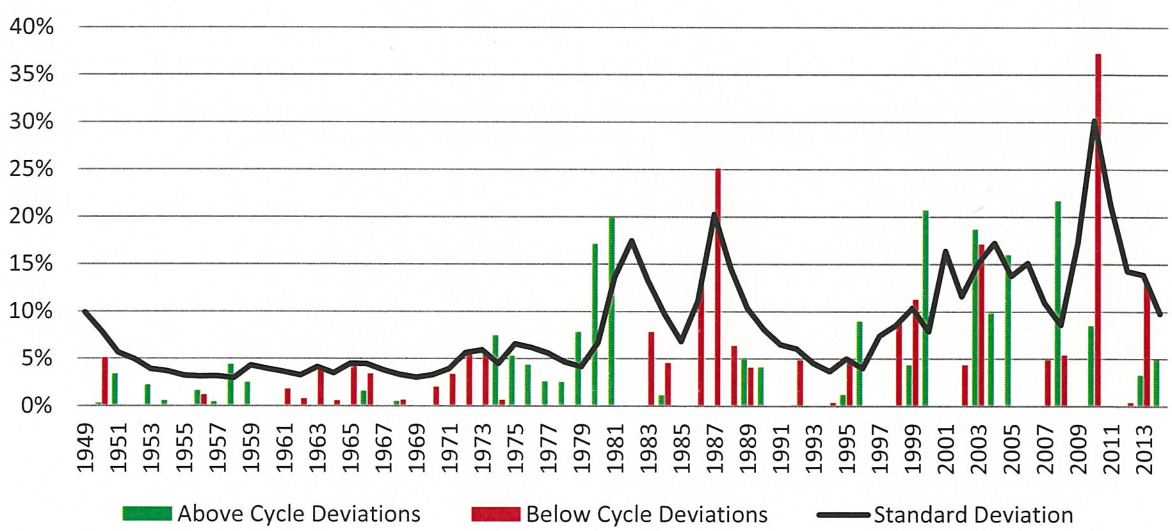

Figure 41 – Characteristics of "Other" Revenues, 1981-2014

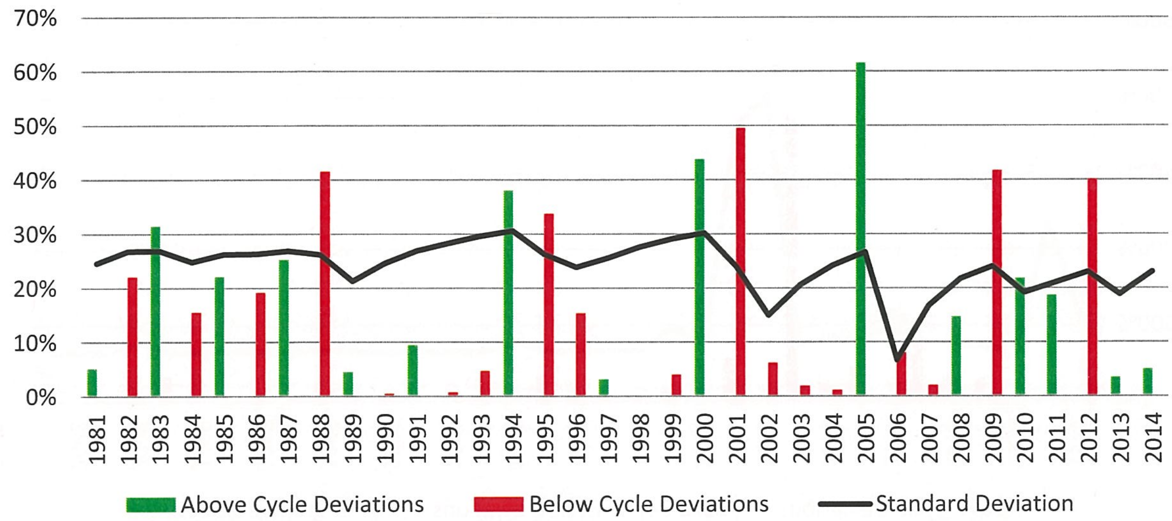

Time-Varying Correlations: A Measure of Volatility Covariation

Figure 42 – Sales and Use Tax Correlations, 1972-2014

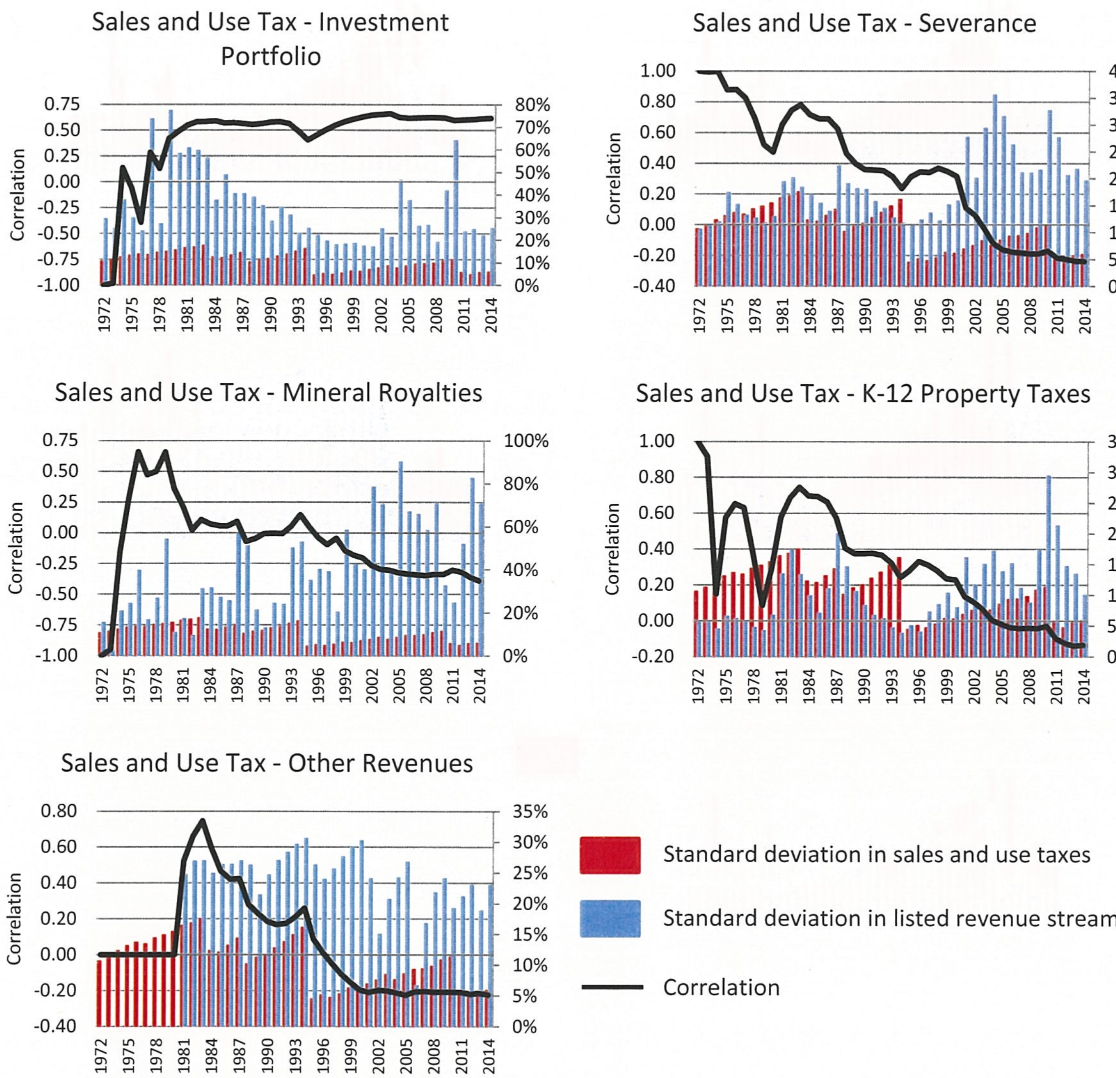

Figure 43 – Investment Income Correlations, 1972-2014

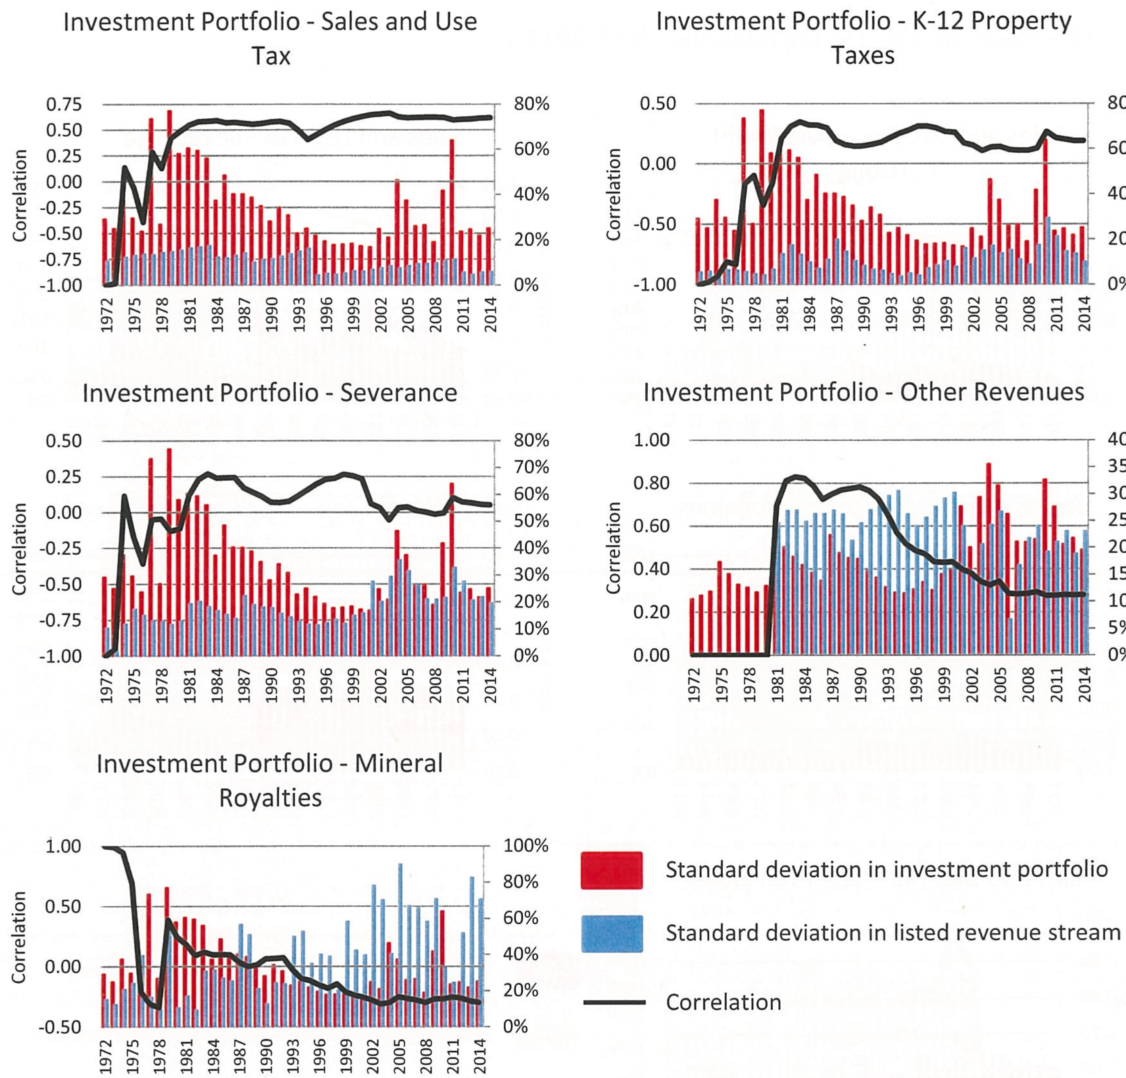

Figure 44 – Severance Correlations, 1972-2014

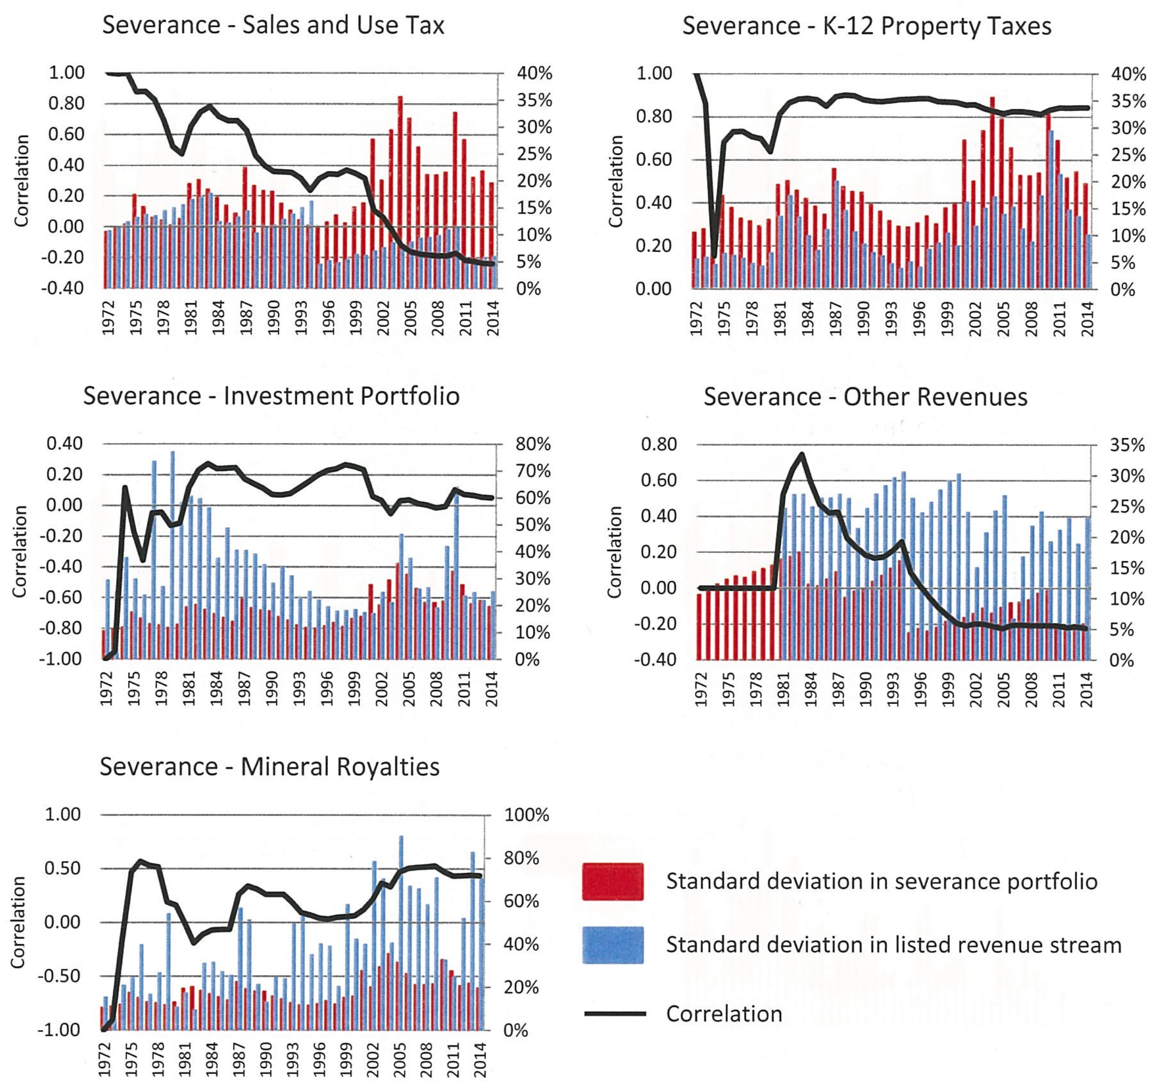

Figure 45 – Mineral Royalties Correlations, 1972-2014

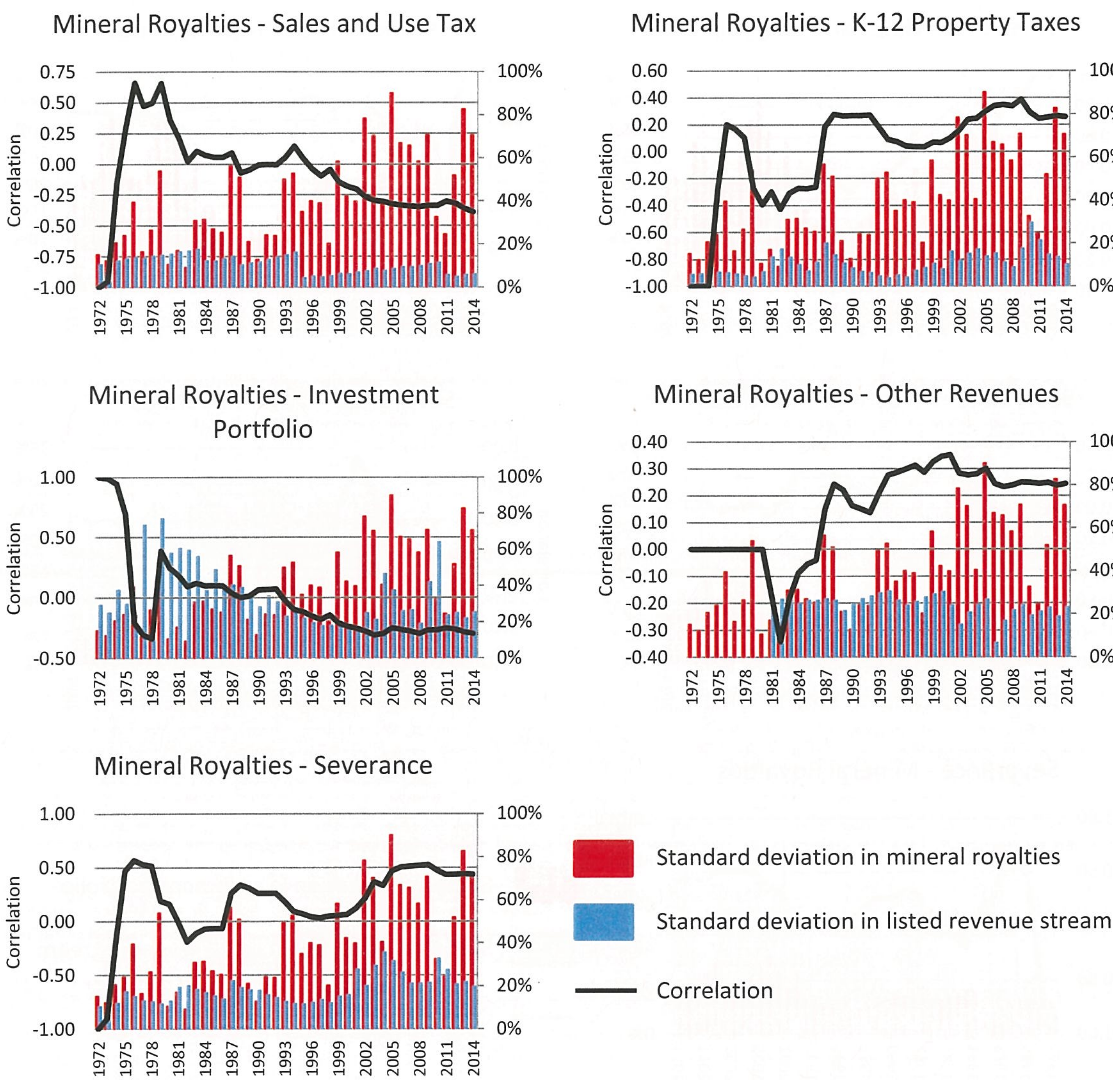

Figure 46 – K-12 Property Tax Correlations, 1972-2014

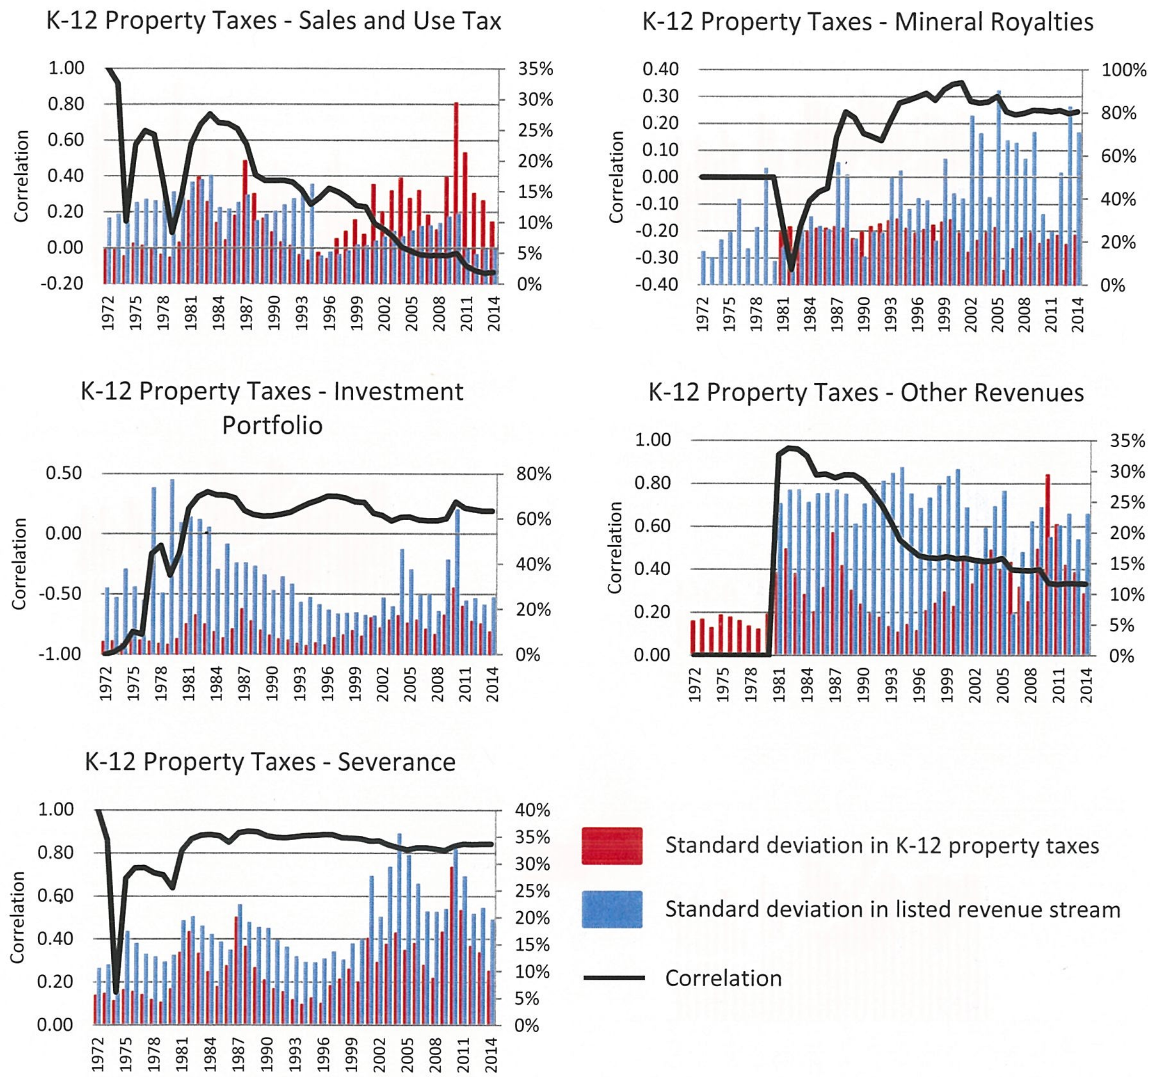

Figure 47 – "Other" Revenue Correlations, 1972-2014

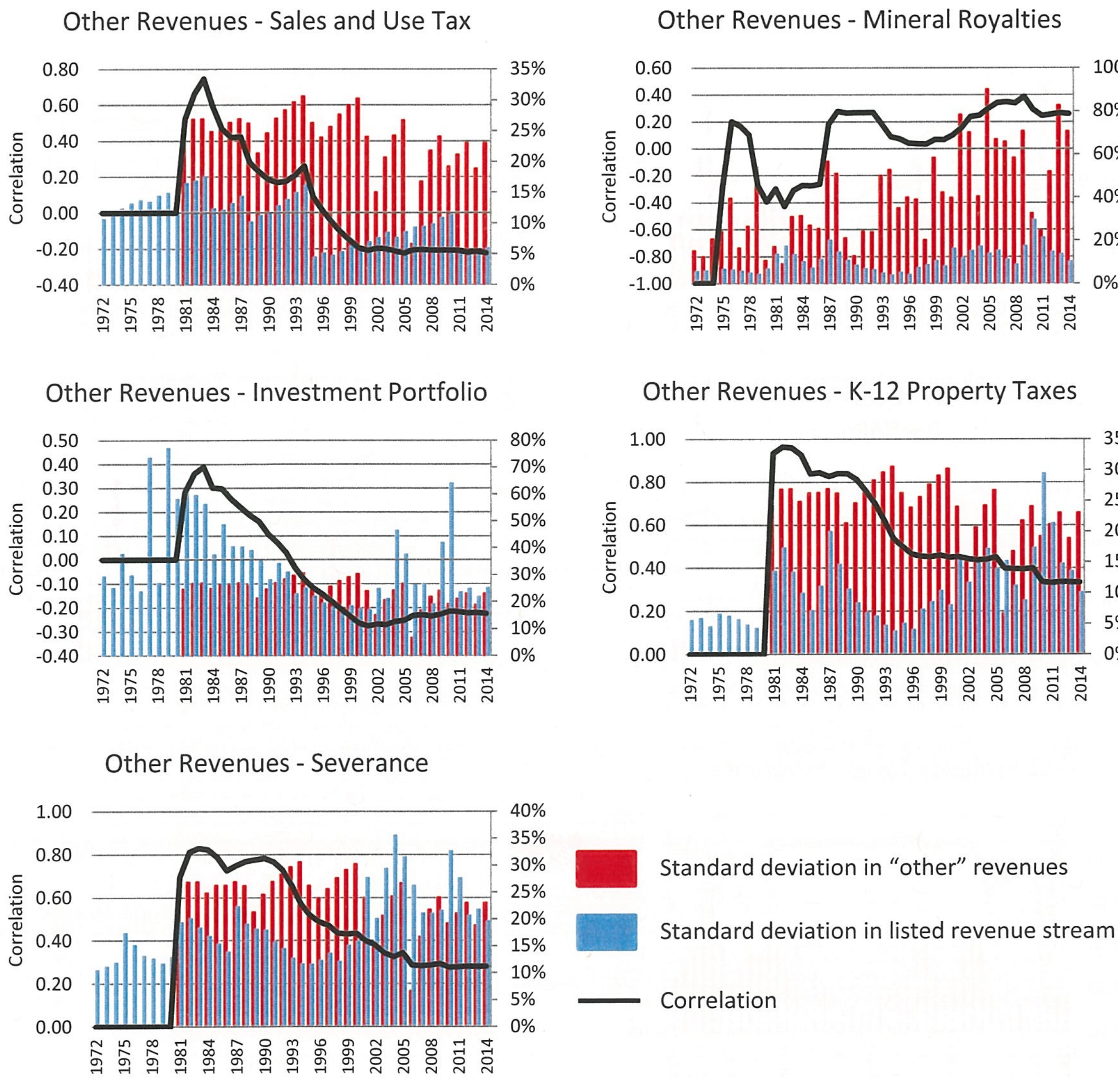

Portfolio Construction: Building a Composite Measure of Volatility

Figure 48 – Investment Portfolio Correlations, 1971-2014

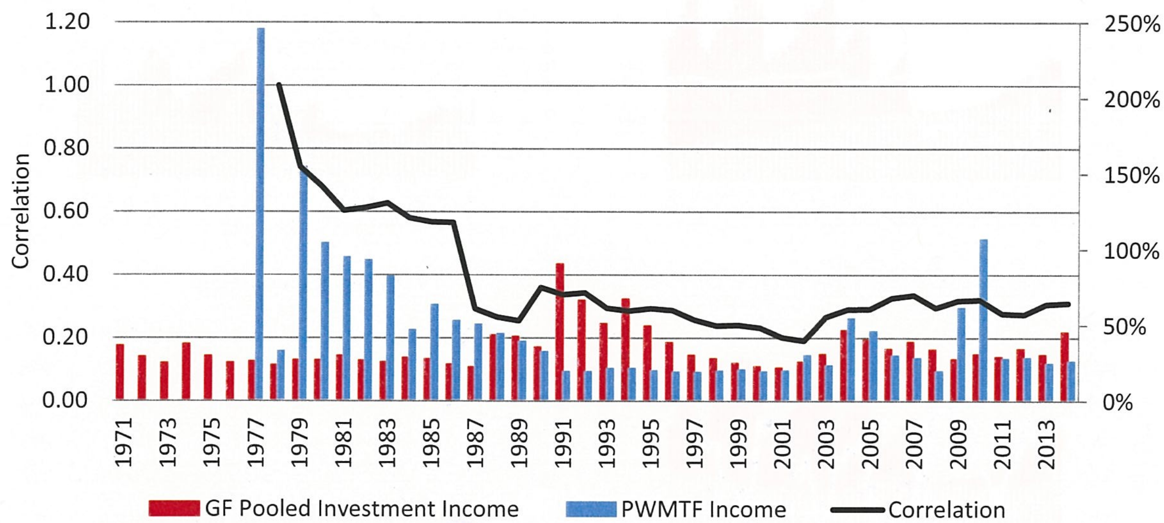

Figure 49 – Investment Portfolio Volatilities, 1971-2014

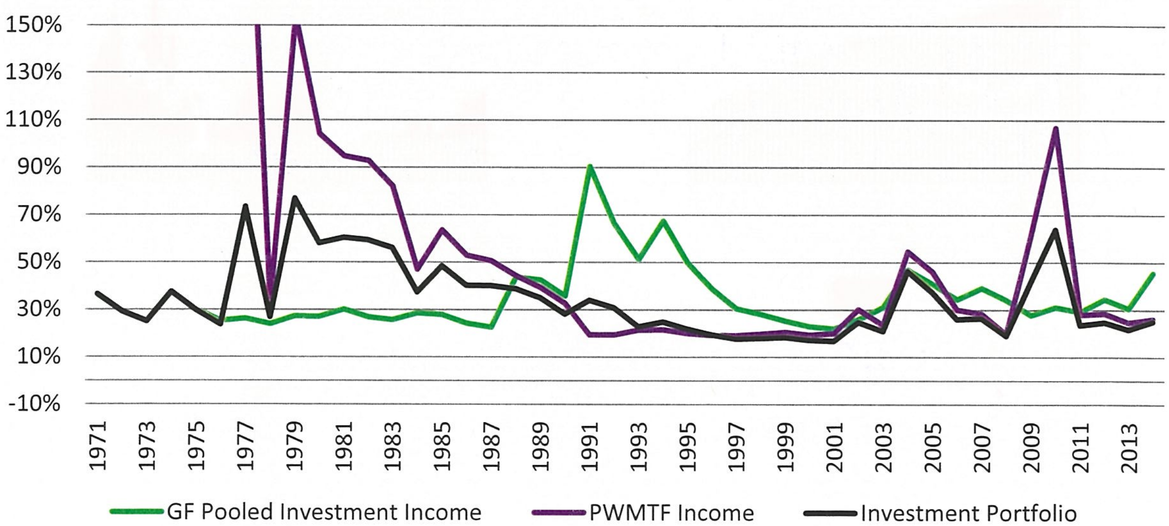

Figure 50 – Severance Portfolio Correlations, 1949-2014

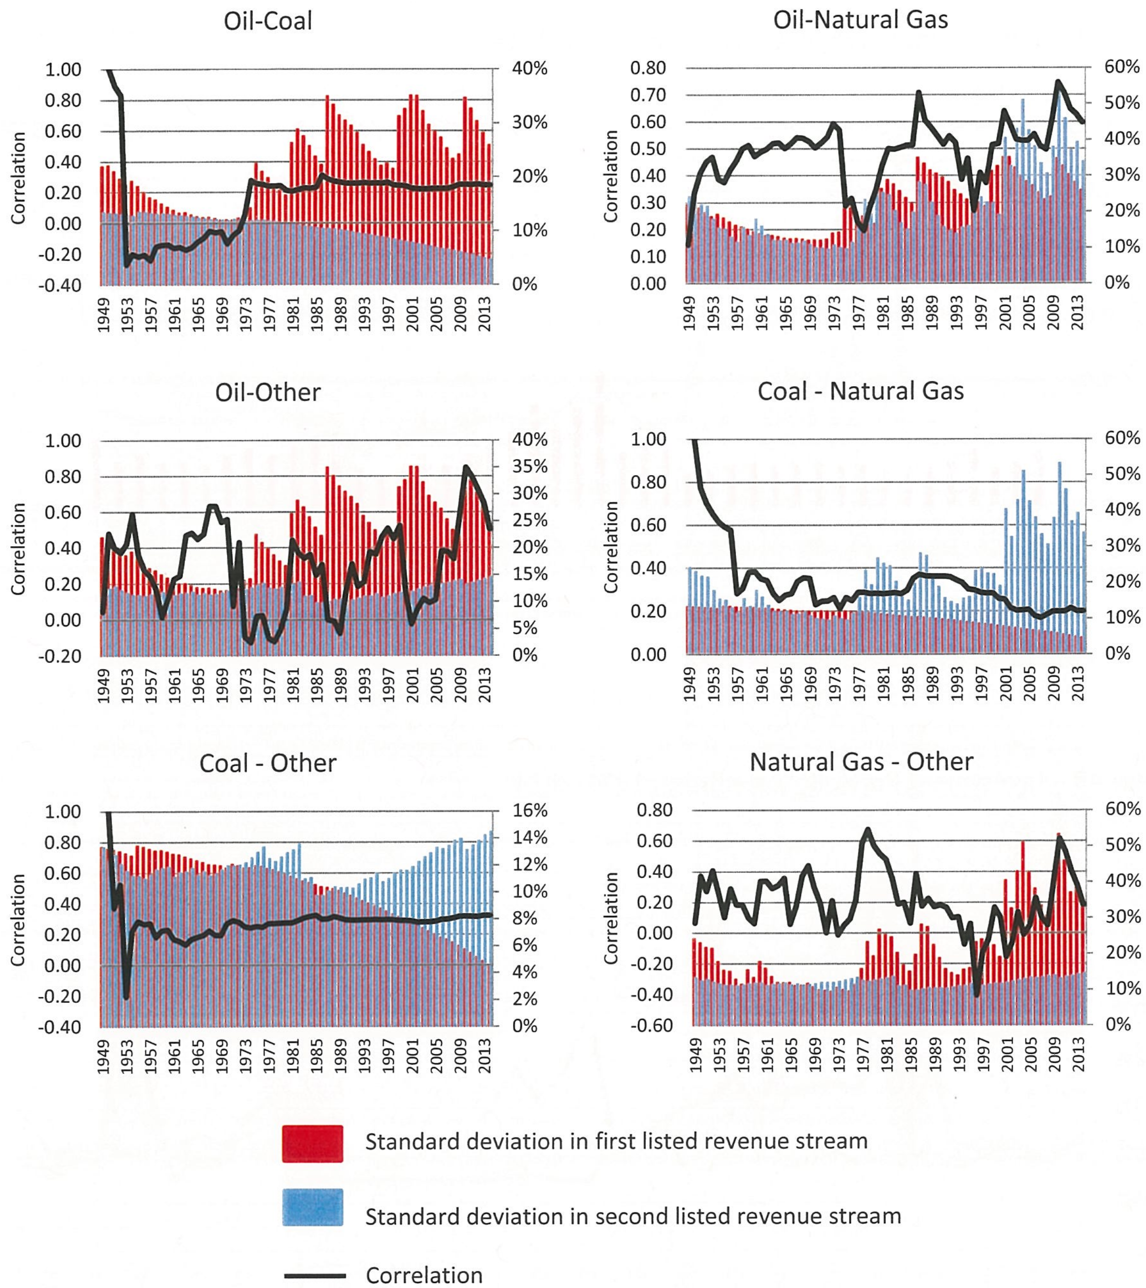

Figure 51 – Severance Portfolio Volatilities, 1949-2014

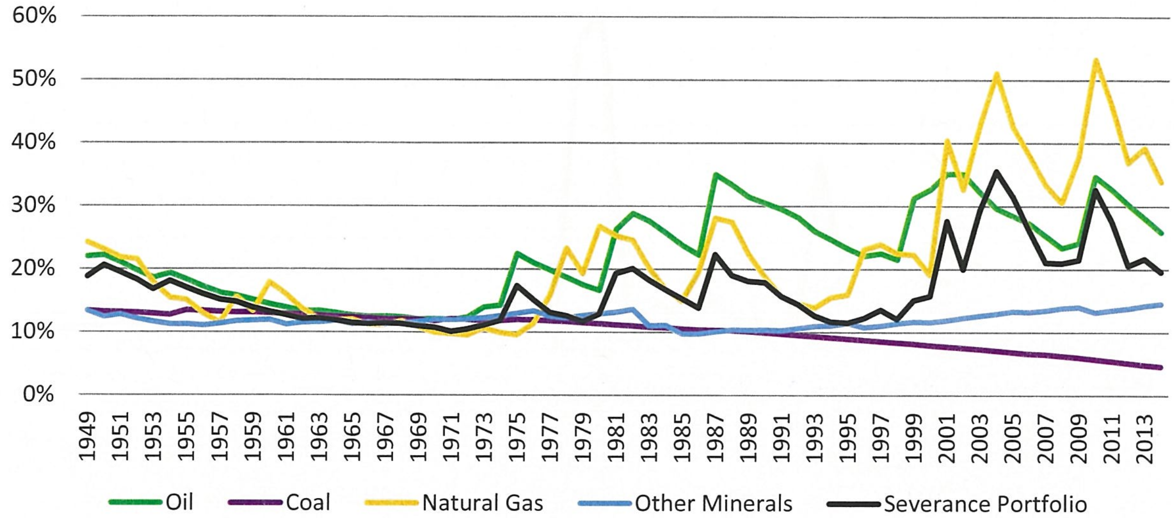

Figure 52 – Mineral Royalties Portfolio Correlations, 1971-2014

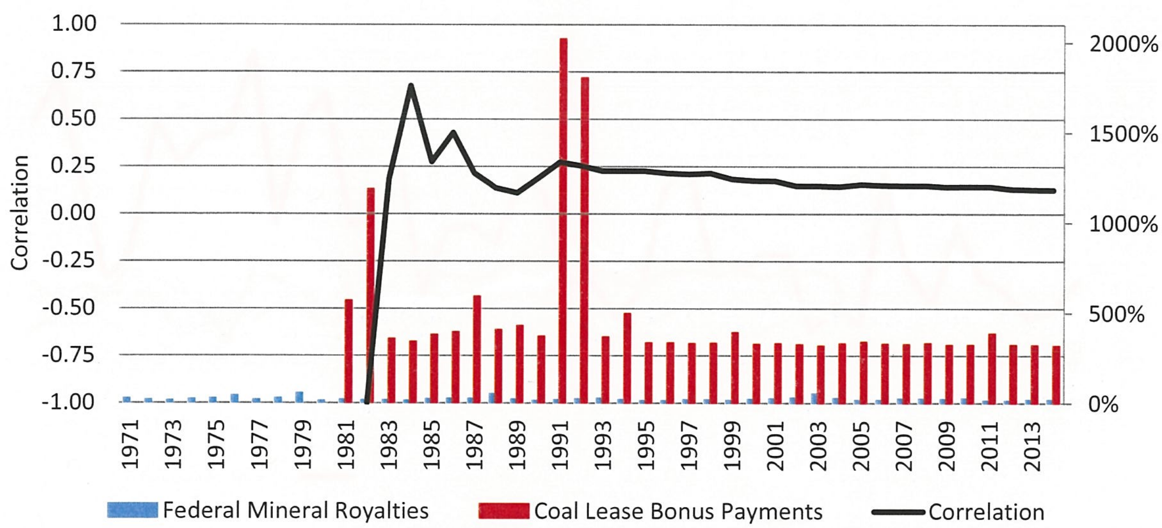

Figure 53 – Mineral Royalties Portfolio Volatilities, 1971-2014

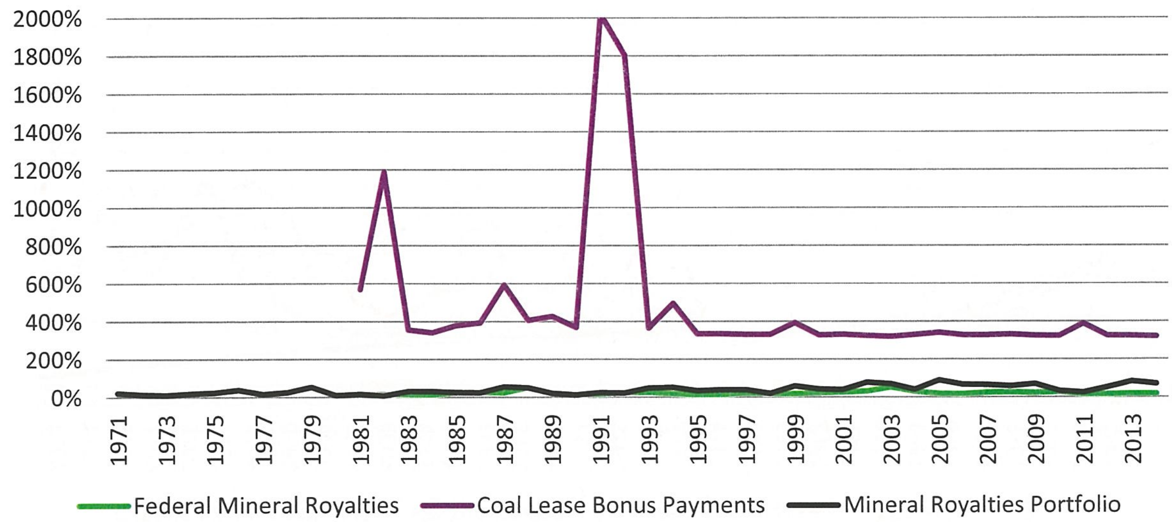

Figure 54 –Full Revenue Portfolio Volatilities, 1971-2014

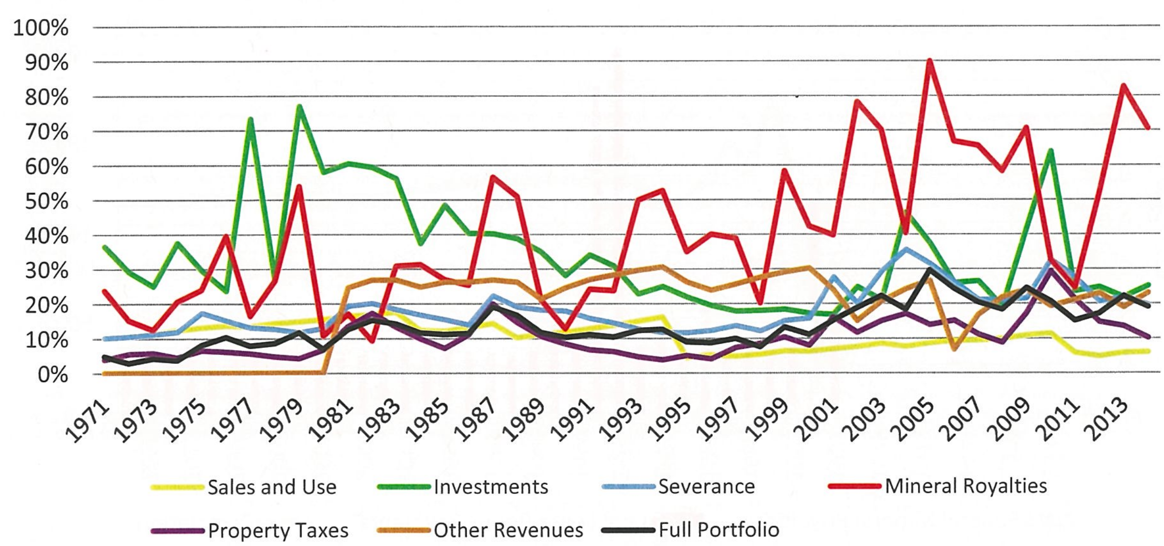

Figure 55 –Full Revenue Portfolio Volatility, 1971-2014

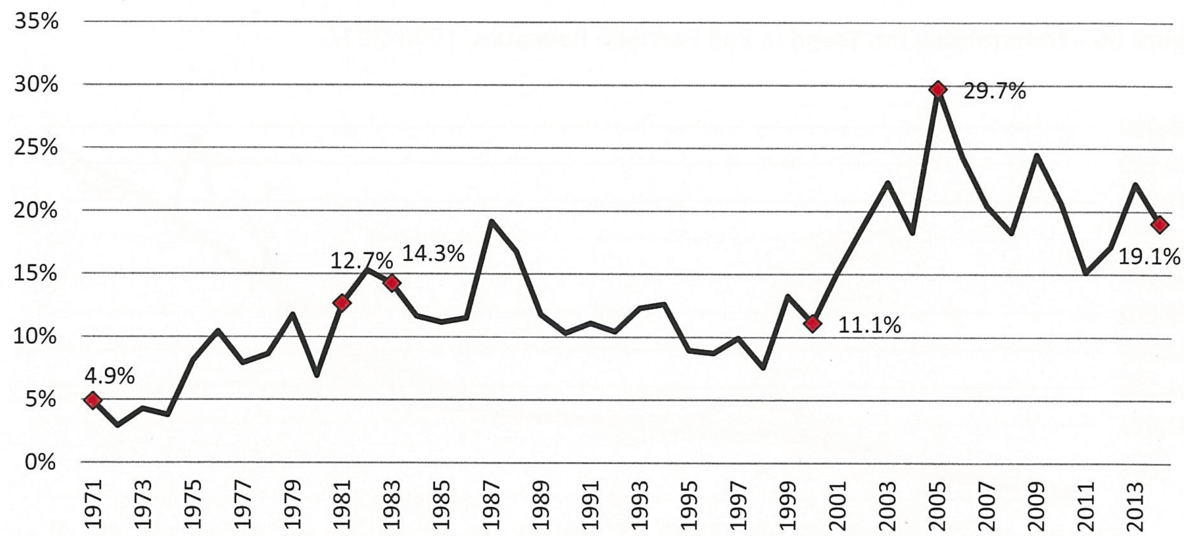

Modeling Persistence in Portfolio Revenues

Figure 56 – Determining the Trend in Full Portfolio Revenues, 1970-2014

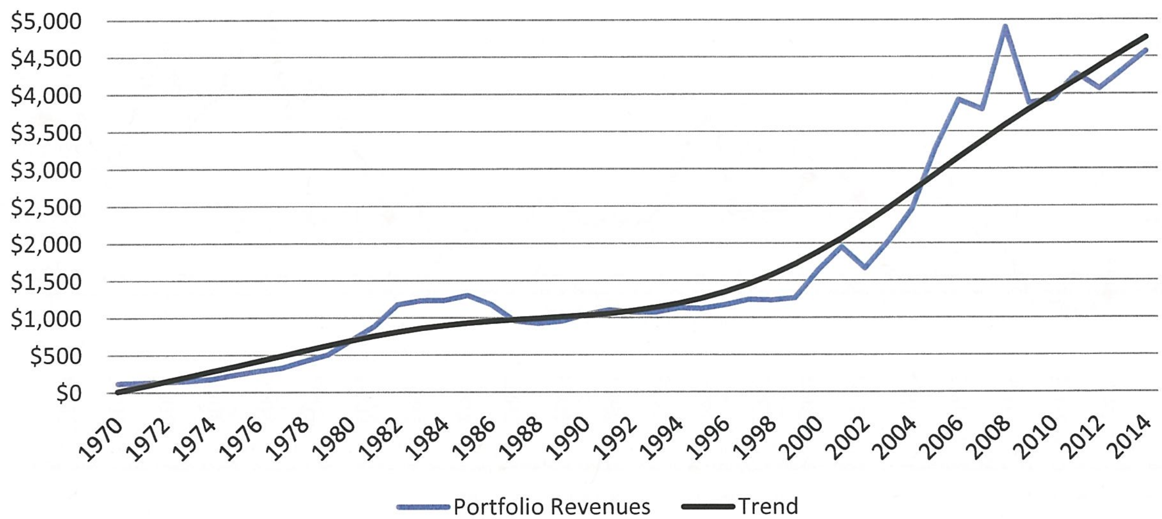

Figure 57 –Full Portfolio Revenues: Deviations from Trend, 1970-2014

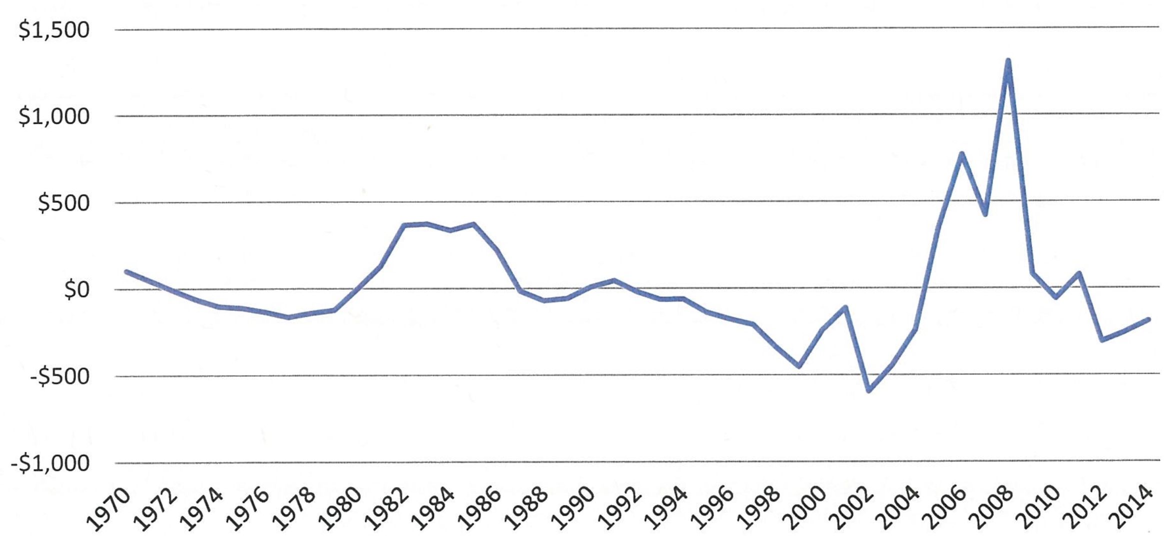

Figure 58 – Explaining the Persistence of Revenue Downturns

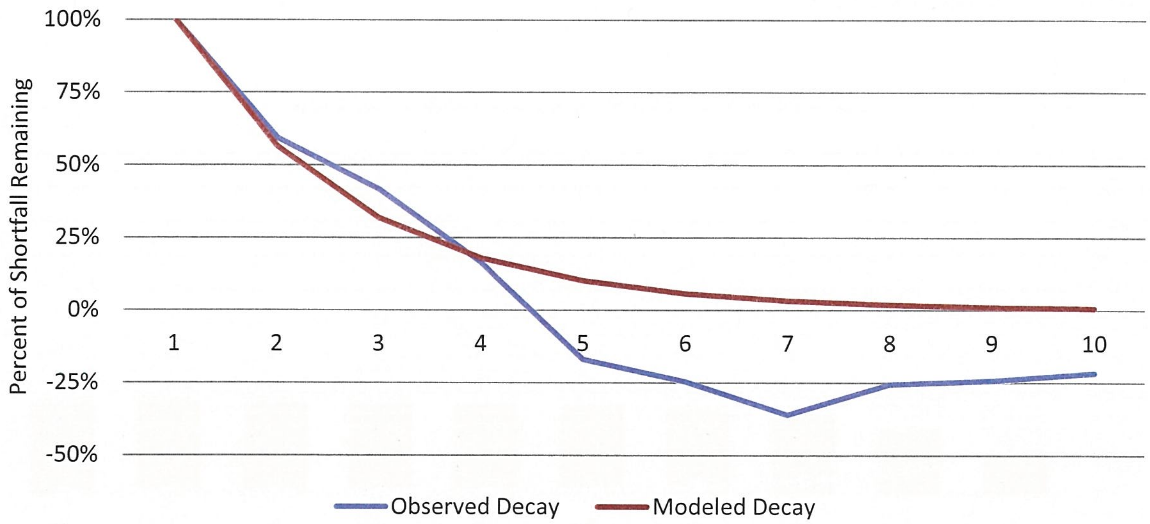

Simulating Revenue Downturns and Savings Needs

Figure 59 – Downturn Scenario 1

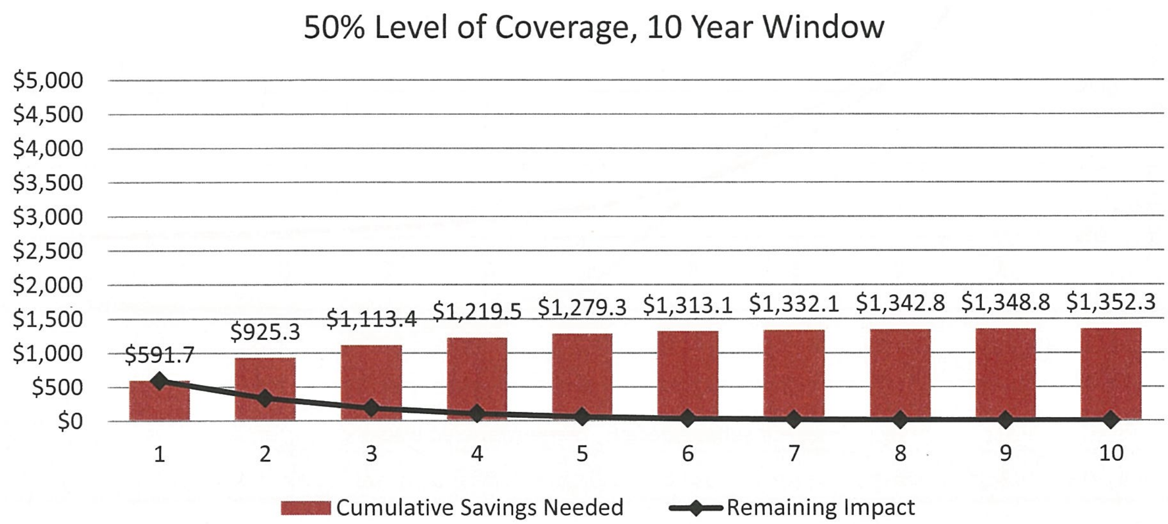

Figure 60 – Downturn Scenario 2

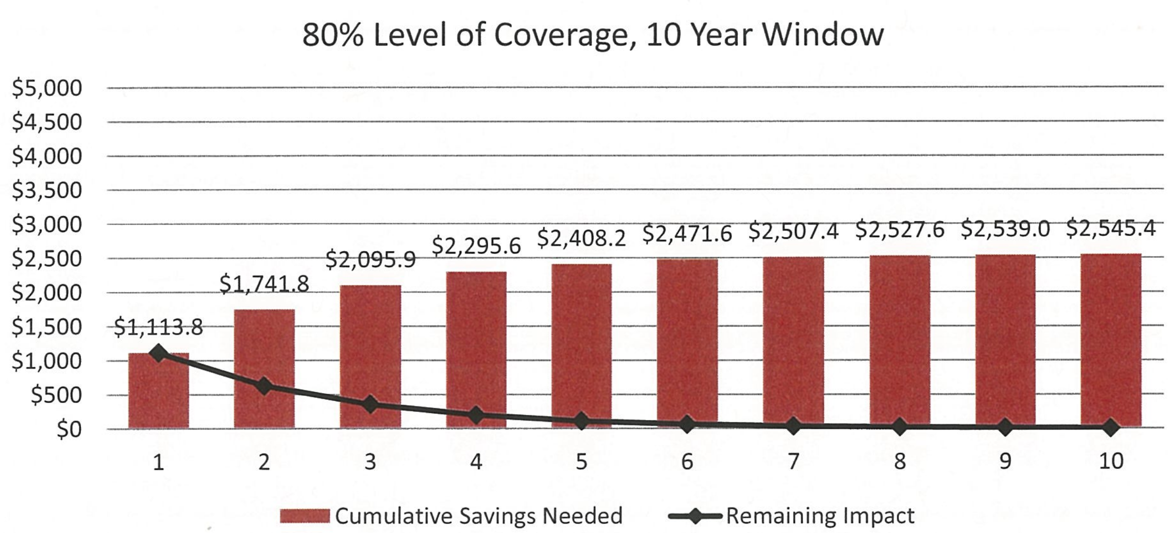

Figure 61 – Downturn Scenario 3