Analysis of Federal Tax Credits Across Income Thresholds Using IRS Filing Season Statistics

Andrew Dymanis, Stat 497 Senior Thesis, Spring 2026

Department of Mathematics and Statistics,
University of Maryland, Baltimore County

Abstract

This paper examines whether average federal tax credit amounts received per return vary systematically across Adjusted Gross Income (AGI) groups using Internal Revenue Service (IRS) Filing Season Statistics. The analysis uses filing season data from 2021 through 2025, which corresponds mainly to tax years 2020 through 2024. Five credits are studied: the Earned Income Credit (EIC), Child Tax Credit (CTC), Refundable Child Tax Credit (RCTC), Education Credit, and Premium Tax Credit (PTC). For each credit and year, average credit per return is computed by dividing total credit amount by the number of returns claiming that credit within each AGI group. The analysis uses descriptive statistics, claimant-count comparisons, graphical comparisons, Pearson correlation, ordinary least squares regression, broad-income-group comparisons, and aggregate credit amount summaries.

The results show that the Earned Income Credit has the clearest income-based structure. It rises through the lower AGI groups, reaches a maximum around the \$20,000 to \$25,000 AGI range, and then declines as income increases. This pattern matches the phase-in and phase-out design of the credit. The Education Credit shows a modest negative relationship with AGI in the correlation analysis. The Child Tax Credit, Refundable Child Tax Credit, and Premium Tax Credit show less direct linear relationships because their amounts depend on family structure, refundability rules, eligibility limits, insurance-market conditions, and the number of claimants within each AGI group. Overall, the results show that federal tax credits are income-sensitive, but the shape, strength, and statistical interpretation of the income relationship differs by credit type.

Key words. Federal tax credits; Adjusted Gross Income; Earned Income Credit; Child Tax Credit; Refundable Child Tax Credit; Premium Tax Credit; IRS Filing Season Statistics; Pearson correlation; linear regression.

AMS subject classifications (2010). 62J05 (Linear regression); 62P20 (Statistics in economics); 91B82 (Statistical methods in economics).

1 Introduction

The federal income tax system is not only a method for collecting revenue. It is also a policy system that distributes financial benefits through deductions, exclusions, and credits. Tax credits are especially important because they directly reduce tax liability. Refundable tax credits can go further by increasing a taxpayer's refund even when the taxpayer has little or no income tax liability. Because of this, credits such as the Earned Income Credit, Refundable Child Tax Credit, and Premium Tax Credit are important tools for supporting low-income and moderate-income households [5–7].

This thesis studies whether average federal tax credit amounts differ across Adjusted Gross Income (AGI) groups. The main research question is whether taxpayers in lower AGI groups receive larger average credit amounts than taxpayers in higher AGI groups. The question matters because many federal credits are not designed to be evenly distributed across income. Some credits have income limits. Some phase in as income rises. Some phase out after a maximum amount is reached. Some depend on the number of children, education expenses, health insurance premiums, or tax liability [9, 10, 12].

The analysis uses Internal Revenue Service (IRS) Filing Season Statistics by AGI group for filing seasons 2021 through 2025. The IRS publishes these tables as part of its Statistics of Income (SOI) program. Filing season statistics summarize tax returns processed at important points during the filing season and report selected income items, adjustments, credits, and taxes by AGI group [1, 2, 4]. The data are useful for this thesis because they give a public and reproducible source for studying how federal tax credits vary across income categories.

The five credits examined are the Earned Income Credit (EIC), Child Tax Credit (CTC), Refundable Child Tax Credit (RCTC), Education Credit, and Premium Tax Credit (PTC). These five credits were selected because they represent different types of federal tax policy. The EIC is a refundable wage-support credit. The CTC is a child-related credit that includes both refundable and non-refundable features. The RCTC represents the refundable portion of child-related relief. The Education Credit is connected to qualified education expenses. The PTC is connected to marketplace health insurance eligibility [5–8, 13].

Table 1.1 lists the main abbreviations used throughout the paper. This table is included so that the tax and statistical terms are clear before the analysis begins.

Table 1.1: Abbreviations used in this thesis.

Abbreviation	Meaning
AGI	Adjusted Gross Income
IRS	Internal Revenue Service
SOI	Statistics of Income
EIC	Earned Income Credit
EITC	Earned Income Tax Credit
CTC	Child Tax Credit
RCTC	Refundable Child Tax Credit
PTC	Premium Tax Credit
OLS	Ordinary Least Squares
ANOVA	Analysis of Variance

The central outcome variable is average credit per return. For each credit type, tax year, and AGI group, average credit per return is computed by dividing the total credit amount by the number of returns claiming that credit. This outcome is useful because it converts total credit dollars into an average amount within each income group. However, the measure must be interpreted carefully. A high average credit in an AGI group does not necessarily mean that all taxpayers in that income group received a large benefit. It means that among returns claiming the credit in that AGI group, the average amount was large.

The expected pattern is not identical across credits. The EIC should show the clearest income relationship because the credit is explicitly designed around earnings and income thresholds [5, 10]. It should rise in the lower income range, reach a maximum, and then decline as income increases. The CTC should not necessarily decline monotonically because taxpayers need qualifying children and, for the non-refundable part of the credit, enough tax liability to use the credit [6, 12]. The RCTC should be concentrated among families with earned income and qualifying children, but the pattern may be uneven in aggregate data. The Education Credit should be smaller and moderately income-sensitive because it depends on qualifying education expenses and income limits [8,11]. The PTC may show the largest total dollar amounts, but it is also likely to be strongly affected by insurance-market eligibility and claimant composition [7, 13].

The statistical goal is not to prove causation at the individual taxpayer level. The IRS data used here are aggregate AGI tables rather than taxpayer-level microdata. Because of this, the thesis estimates and interprets group-level relationships. The analysis is still useful because it shows whether public IRS summary tables contain visible distributional patterns consistent with the structure of federal tax credits.

The thesis proceeds as follows. Section 2 explains the credit designs and expected income gradients. Section 3 describes the data source, table structure, variables, AGI group

sizes, and claimant counts. Section 4 defines the statistical methods. Section 5 presents the descriptive results, figures, correlation results, regression results, and broad income group comparisons. Section 6 summarizes the findings, policy implications, limitations, and extensions.

2 Credit Design and Expected Income Gradients

The federal credits studied in this thesis have different eligibility rules. Because of this, AGI should not affect each credit in the same way. A correct statistical interpretation must compare the data against the design of each credit. This section explains the expected pattern for each credit before the empirical results are presented.

The **Earned Income Credit (EIC)** is a refundable credit designed for low-income and moderate-income workers. The credit generally phases in as earned income increases, reaches a maximum, and then phases out as income rises further [5, 10]. This structure creates a nonlinear relationship between AGI and credit amount. The EIC is not expected to be highest at the very lowest or very highest income levels. Instead, it should peak in a lower-to-moderate income range and then decline. This is important for the statistical analysis because a simple Pearson correlation or straight-line regression may not fully capture the EIC pattern.

The **Child Tax Credit (CTC)** provides a credit for qualifying children. It is different from the EIC because the credit is not only a wage support program. It depends on the number of qualifying children and, for the non-refundable portion, whether the taxpayer has enough tax liability to use the credit [6,12]. Because of this, the CTC can be larger in middle-income groups than in the lowest AGI groups. Taxpayers with very low income may have children but not enough tax liability to receive the full non-refundable amount. Therefore, a positive relationship between AGI and average CTC in some ranges is not automatically inconsistent with the credit's design.

The **Refundable Child Tax Credit (RCTC)** is the refundable portion of the child-related credit system. Its value depends on qualifying children, earned income, and refundability rules [6, 12]. The expected relationship with AGI is complicated. A low-income household with children may receive a substantial refundable amount, but a household with no earned income or no qualifying children may receive little or nothing. In aggregate AGI data, this can produce spikes in certain income brackets. Those spikes may reflect the composition of claimants rather than a smooth income gradient.

The **Education Credit** includes education-related credits such as the American Opportunity Credit and Lifetime Learning Credit. These credits are tied to qualified education expenses and income limits [8,11]. They are not expected to dominate the total federal credit distribution because only taxpayers with qualifying education expenses can claim them. The

expected relationship with AGI is moderately negative because income limits reduce eligibility at higher levels, but the effect should be weaker than the EIC because education expense eligibility is not directly tied to low-income work in the same way.

The **Premium Tax Credit (PTC)** subsidizes marketplace health insurance premiums for eligible households. Its value depends on household income, insurance costs, household size, state-level insurance prices, and marketplace eligibility [7, 13]. Therefore, the relationship between AGI and the PTC may not appear as a simple negative slope in aggregated data. A higher AGI group can show a high average PTC if the returns that claim the credit in that bracket have high eligible premiums. This means the PTC should be interpreted using both average credit per return and total credit amount.

The key point is that AGI is an important predictor, but it is not the only determinant of credit amount. The EIC should show the cleanest income gradient because income directly drives the statutory formula. The CTC and RCTC depend heavily on family structure. The Education Credit depends on education expenses. The PTC depends on insurance eligibility and premium costs. These differences explain why the empirical results are expected to vary by credit.

3 Data

3.1 Source and Background

The data come from the IRS Filing Season Statistics page [1]. The IRS publishes this page as a public statistical resource summarizing return filing activity during each filing season. The tables report selected information from individual income tax returns, including numbers of returns, income items, adjustments, deductions, credits, and taxes. The IRS also publishes related SOI tables by AGI group and filing status, which provide background for how tax data are commonly organized in public IRS statistical products [2–4].

The files used in this thesis were found by using the IRS Filing Season Statistics page and selecting the individual return tables by filing season and cycle. Each file contains a table titled *All Individual Returns: Selected Income Items, Adjustments, Credits, and Taxes, by Size of Adjusted Gross Income*. This table is appropriate for the thesis because it directly reports tax credit information by AGI group.

The analysis uses filing seasons 2021 through 2025. These filing seasons correspond mainly to tax years 2020 through 2024 because individual income tax returns are generally filed in the year after the tax year. The year labels in the graphs are kept as 2021 through 2025 because those are the filing season years used in the project files and figures.

3.2 Raw File Structure

The raw files are Excel workbooks. Each workbook contains one main worksheet named *Filing Season*. The number of rows and columns varies slightly by year because the IRS table layout changed over time. Table 3.1 summarizes the raw table size for each filing season file used in the analysis.

Table 3.1: Raw IRS filing season files used in the analysis.

Filing Season	Approximate Tax Year	Worksheet	Raw Table Size
2021	2020	Filing Season	103 rows by 19 columns
2022	2021	Filing Season	155 rows by 19 columns
2023	2022	Filing Season	143 rows by 19 columns
2024	2023	Filing Season	141 rows by 18 columns
2025	2024	Filing Season	137 rows by 18 columns

The rows in each file describe tax items, such as all returns, adjusted gross income, salaries and wages, tax credits, and tax amounts. The columns describe AGI groups, such as no adjusted gross income or deficit, $1 under $5,000, $5,000 under $10,000, and higher AGI ranges. Because the data are arranged in wide table form, the data were cleaned and converted into a long format for statistical analysis. The long format has one row for each filing season, AGI group, and credit type.

3.3 Sample of the Raw Data

Table 3.2 shows a small sample of the raw IRS table structure. The purpose of this table is to make the data source more concrete. The full raw files contain many more rows and AGI columns, but this sample shows how the information is organized before the data are cleaned.

Table 3.2: Sample rows from the IRS filing season table structure.

Item	Total Returns	No AGI/Deficit	$1–$5k	$5k–$10k
All returns	140,210,311	1,769,986	5,692,363	6,500,898
Electronically filed returns	137,025,974	1,561,489	5,385,767	6,264,899
Adjusted gross income	12,121,228,502	-48,106,585	13,944,109	49,081,925
Total income: number of returns	138,484,830	44,576	5,692,292	6,500,898
Total income: amount	12,213,243,003	-52,990,196	14,537,713	50,813,349

The sample shows why cleaning is necessary. The original IRS file is not organized as

one observation per row. Instead, tax items are listed down the rows, and AGI groups are listed across the columns. To analyze credits by income group, the credit rows must be selected, the AGI columns must be gathered into one variable, and the credit amounts and claimant counts must be matched. The data cleaning and plotting were done in R, using the R statistical environment and tidyverse workflow [19, 20].

3.4 Description of AGI Group Sizes and Claimant Counts

The income ranges are not equally large. Some AGI brackets are narrow, such as $1 under $5,000, while others are much wider, such as $100,000 under $200,000 or $200,000 under $500,000. The number of returns in each bracket is also not equal. For example, in the final filing season file, the $100,000 under $200,000 group contains more than 25 million returns, while the $1,000,000 or more group contains fewer than one million returns. This matters because average credit values can be affected by both the width of the income bracket and the number of returns or claimants in that bracket.

Table 3.3 gives an example of the number of returns by selected AGI groups for the final year of data. The purpose is to show that the AGI groups are not the same size and should not be interpreted as equally weighted income categories.

Table 3.3: Example AGI group sizes from the final IRS file.

AGI Group	Number of Returns
No adjusted gross income and deficit	1,769,986
$1 under $5,000	5,692,363
$5,000 under $10,000	6,500,898
$10,000 under $15,000	7,598,844
$15,000 under $20,000	7,486,368
$20,000 under $25,000	6,934,913
$50,000 under $75,000	21,799,485
$75,000 under $100,000	14,090,140
$100,000 under $200,000	25,051,983
$1,000,000 or more	403,536

The professor's feedback also asked for claimant counts by income range. Table 3.4 gives the total number of credit-claiming returns by broad income group and credit type across the available filing seasons. This table shows that the number of claimants is not evenly distributed across income groups or across credits. For example, the EIC has no high-income claimants under the broad grouping used here, while the PTC and RCTC have claimants in all three broad income groups.

Table 3.4: Credit claimant counts by broad income group across the filing season data.

Credit Type	Low Income	Middle Income	High Income
Child Tax Credit	12,997,234	17,739,767	2,611,118
Earned Income Credit	111,702,617	3,180,510	0
Education Credit	50,701,594	43,774,580	0
Premium Tax Credit	244,226,504	90,193,412	3,009,820
Refundable Child Tax Credit	199,424,211	169,964,102	27,821,523

Table 3.4 is important because average credit amounts are based only on returns claiming each credit. The low-income group has the largest number of EIC claimants, which matches the credit's purpose. The PTC and RCTC have much larger total claimant counts because those credits appear in many AGI groups. The CTC is shown only for the available filing season in the cleaned credit table, which is why its total claimant count is lower than the multi-year counts for some other credits.

Claimant counts also differ by credit. A credit can have a high total amount because many taxpayers claim it, because the average benefit is large, or because both are true. For that reason, the analysis uses both average credit per return and total credit amount. The average credit measure answers how large the credit is among claimants in an AGI group. The total amount measure answers how large the credit is in aggregate dollars.

3.5 Variable Construction

The main outcome variable is average credit per return:

$$\bar{c}_{g,t,k} = \frac{\text{Total credit amount}_{g,t,k}}{\text{Number of returns claiming credit}_{g,t,k}}, \tag{3.1}$$

where g indexes AGI group, t indexes filing season year, and k indexes credit type.

This measure is interpreted as the average amount received among returns claiming the credit in that AGI group. It is not the same as average credit per all taxpayers in the bracket. This distinction is important. If a bracket has a small number of claimants, the average credit among those claimants can be large even if the credit is not common in that bracket.

Because the IRS reports income in brackets, each AGI bracket is assigned a midpoint. These midpoint values are used as the explanatory variable in the correlation and regression analysis. For example, the bracket \$20,000 under \$25,000 is represented by a midpoint near \$22,500. For very high brackets, the midpoint approximation becomes less precise because the bracket width is larger and the distribution of income inside the bracket may be skewed.

For broad income group comparisons, AGI groups are classified as:

$$\text{income group}_g = \begin{cases} \text{Low income} & m_g < 50{,}000, \\ \text{Middle income} & 50{,}000 \leq m_g < 200{,}000, \\ \text{High income} & m_g \geq 200{,}000. \end{cases} \tag{3.2}$$

This classification is used to compare average credits across low, middle, and high income groups. It also makes the results easier to interpret because broad income categories are more understandable than a long list of AGI brackets.

3.6 Analytical Unit

The analytical unit is the credit-year-AGI group cell. For example, one cell could represent the average EIC among returns claiming the EIC in the $20,000 under $25,000 AGI group in filing season 2025. Another cell could represent the average PTC among returns claiming the PTC in the $100,000 under $200,000 AGI group in filing season 2024.

This structure means the analysis is ecological rather than individual-level. The results describe patterns across AGI groups, not causal effects for individual taxpayers. This is appropriate for the research question because the thesis asks whether the public IRS AGI tables reveal income-based credit patterns.

3.7 Important Measurement Issue

A major measurement issue is the difference between average credit per claiming return and average credit per tax return overall. The graphs use average credit per return based on credit-specific claimant counts. This makes the average meaningful for returns that claimed the credit, but it also means that the denominator changes by credit type, year, and AGI group.

For the EIC, this is less problematic because the credit is closely tied to income eligibility and claimant patterns are part of the policy structure. For the PTC and RCTC, interpretation is harder because a small number of claimants in a high-income group can create a large average. This is one reason the thesis uses multiple views of the data: average credit graphs, total credit amounts, correlations, regression lines, and broad income group comparisons.

4 Methodology

4.1 Descriptive Statistics and Graphs

The first part of the analysis is descriptive. The analysis begins by checking the IRS file structure, identifying the AGI groups, and summarizing the number of returns and credit claimants across income ranges. Descriptive statistics are necessary because the income ranges are not equally large and because the number of claimants can differ substantially by credit type.

The graphical analysis shows average credit per return by AGI group, total credit amount by year, average credit by broad income group, correlation heatmaps, and regression lines. These graphs are used to identify whether each credit has a clear income pattern.

Graphical analysis is necessary because not all credit relationships are linear. The EIC is the clearest example. Because the credit phases in and then phases out, a graph gives more information than a single correlation coefficient. A single statistic can hide the policy shape of the credit.

4.2 Pearson Correlation

For each credit-year pair, the Pearson correlation is computed between AGI midpoint and average credit per return:

$$r_{k,t} = \frac{\sum_{g=1}^{G}(m_g - \bar{m})(\bar{c}_{g,t,k} - \bar{c}_{t,k})}{\sqrt{\sum_{g=1}^{G}(m_g - \bar{m})^2 \sum_{g=1}^{G}(\bar{c}_{g,t,k} - \bar{c}_{t,k})^2}}. \tag{4.1}$$

Pearson correlation measures the strength and direction of a linear association between two quantitative variables [14, 15]. A negative value means that average credit per return tends to decline as AGI rises. A positive value means that average credit per return tends to increase as AGI rises. A value close to zero means the relationship is weak, nonlinear, or affected by outliers.

Pearson correlation is useful as a first diagnostic, but it is not sufficient by itself. It assumes the main pattern can be summarized by a straight-line association. For credits with nonlinear statutory rules, this can be misleading. Therefore, the correlation heatmap is interpreted alongside the line graphs and regression plots.

4.3 Simple Linear Regression

For each credit-year cell, the following simple ordinary least squares (OLS) regression model is estimated:

$$\bar{c}_{g,t,k} = \beta_{0,k,t} + \beta_{1,k,t} m_g + \varepsilon_{g,t,k}. \tag{4.2}$$

The slope $\beta_{1,k,t}$ measures the average change in credit per return associated with a one-dollar increase in AGI midpoint. A negative slope supports the hypothesis that average credit declines as income rises. A positive slope means the average credit is higher in larger AGI groups. Linear regression is a standard method for summarizing the average linear relationship between a response variable and one or more explanatory variables [15, 16, 18].

The regression model is useful because it provides a simple numerical summary of direction. However, it has an important limitation. Some credits, especially the EIC, are nonlinear by design. The EIC first rises, then falls. A straight-line regression can understate this type of relationship. For this reason, the regression plots are interpreted as descriptive summaries, not as complete models of credit formulas.

4.4 Analysis of Variance

Analysis of variance (ANOVA) is used as an additional way to compare average credit values across broad income groups. ANOVA is a standard statistical method for comparing group means [14, 17]. In this thesis, the broad groups are low income, middle income, and high income as defined in Equation 3.2.

The ANOVA interpretation is cautious because the observations are aggregate AGI cells rather than individual tax returns. Therefore, ANOVA is used as descriptive evidence about group differences, not as a causal test of how income changes an individual taxpayer's credit.

4.5 Interpretation Strategy

The interpretation follows three rules. First, each credit is interpreted according to its design. The EIC is expected to show a phase-in and phase-out pattern, while the CTC is expected to depend more on family structure and tax liability. Second, no single statistic is treated as decisive. The thesis compares graphs, correlations, regression lines, group comparisons, claimant counts, and aggregate totals. Third, high average values in upper AGI groups are interpreted cautiously because they may reflect a small number of claimants rather than a broad distribution of benefits.

This strategy is important because aggregated IRS data can show real patterns but can also produce misleading averages when cell counts are small or when eligibility depends on unobserved household characteristics.

5 Results

5.1 Descriptive Statistics

The results begin with a concrete description of the data because the AGI groups are not equally sized and the number of returns differs substantially across income ranges. The raw IRS files contain between 103 and 155 rows and between 18 and 19 columns, depending on the year. After cleaning, the analysis focuses on five credit categories across five filing seasons and the AGI groups available in each IRS table.

The claimant-count summary in Table 3.4 shows that credit participation is not evenly distributed across income groups. This matters because the average credit per return is calculated among claiming returns only. A credit with more claimants in one income group can have a much different aggregate effect than a credit with only a few claimants in that same group.

Table 5.1 summarizes the main interpretation of each credit in the data. The table is qualitative because the exact values vary by year and AGI group.

Table 5.1: Summary interpretation of credit patterns in the IRS AGI data.

Credit Type	Main Pattern in the Data
Earned Income Credit	Strong nonlinear pattern. The average credit rises in low AGI groups, peaks near the \$20,000 to \$25,000 range, and then declines as AGI increases.
Child Tax Credit	Increases across lower and middle AGI groups in the year shown. This reflects that the credit is not purely low-income targeted and depends on qualifying children and tax liability.
Refundable Child Tax Credit	Uneven pattern. Some years show large values in specific AGI groups, indicating that family structure and refundability rules strongly affect the average credit.
Education Credit	Smaller average credit values overall. The correlation is consistently negative in the years shown, but the magnitude is modest.
Premium Tax Credit	Largest aggregate dollar amount. Average values vary sharply across AGI groups, especially in higher AGI groups, suggesting strong compositional effects in the aggregated data.

The descriptive results show that the credits should not be interpreted as if they all follow the same statistical pattern. The EIC is the most directly income-sensitive credit. The PTC

is the largest by total dollars. The CTC and RCTC are strongly affected by family and refundability rules.

5.2 Graphical Overview of Average Credit by AGI Group

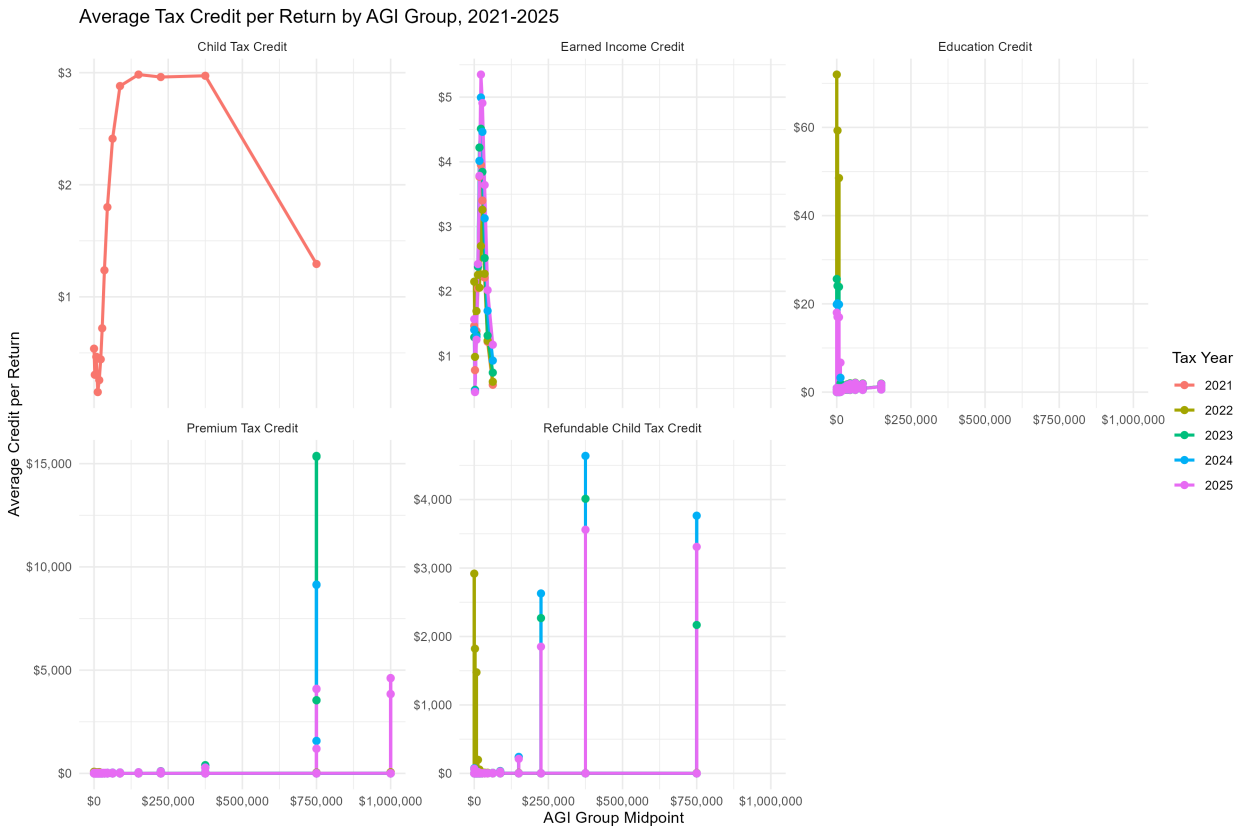


Figure 5.1: Average tax credit per return by AGI group midpoint for the Child Tax Credit, Earned Income Credit, Education Credit, Premium Tax Credit, and Refundable Child Tax Credit.

Figure 5.1 gives the broadest view of the data. The figure compares average credit per return across AGI group midpoints for all five credits. The main result is that the credit-income relationship differs by credit type.

The EIC has the clearest and most interpretable pattern. Average EIC rises through the lower AGI groups, reaches a maximum in the lower-to-middle AGI range, and then declines as AGI increases. This matches the design of the credit, which phases in, reaches a maximum, and then phases out [5, 10]. The EIC pattern is important because it shows that the relationship between AGI and credit amount is not always linear.

The CTC behaves differently. Average CTC rises from the lowest AGI groups into the middle AGI groups before declining at higher values. This is consistent with the credit being partly non-refundable and dependent on qualifying children [6, 12]. A low-income taxpayer may have children but not enough tax liability to receive the full non-refundable amount. A middle-income taxpayer with qualifying children may receive more. Therefore, the upward pattern in the CTC should not be interpreted as an error. It reflects the credit's design and claimant composition.

The Education Credit has smaller average values for most AGI groups, although some low-AGI points are noticeably higher. The PTC and RCTC have larger spikes in some AGI groups. These spikes show that average values can be sensitive to the number and type of claimants in each bracket. For these credits, high average values in certain brackets may reflect claimant composition rather than a smooth income relationship.

5.3 Earned Income Credit Results

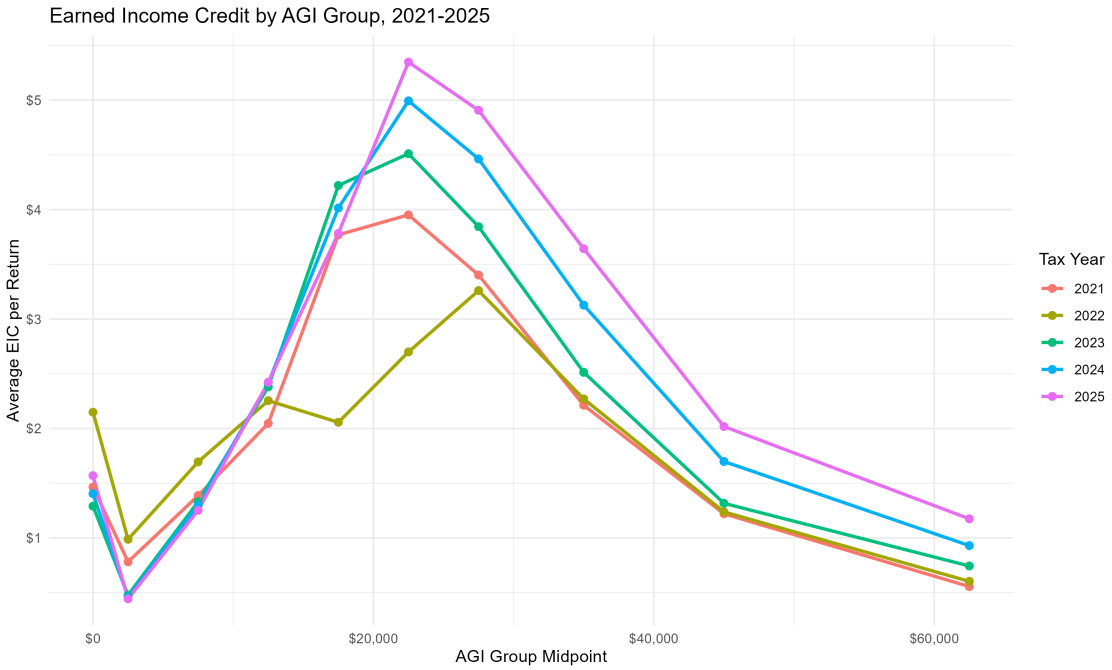


Figure 5.2: Average Earned Income Credit per return by AGI group.

Figure 5.2 isolates the EIC. This is the most important graph for the central thesis question because the EIC is the clearest refundable credit targeted toward lower-income workers.

Across all years, the EIC follows a nonlinear pattern. The average credit is relatively low at the very lowest AGI group, rises quickly through the lower-income brackets, peaks around the $20,000 to $25,000 AGI range, and then declines steadily. This graph supports the thesis

because it shows that the EIC is not randomly distributed across AGI groups. The highest values are concentrated in the lower-to-middle income range. The decline after the peak is the visual evidence of the phase-out effect.

The EIC also explains why correlation alone can be misleading. A straight-line correlation asks whether higher AGI is associated with higher or lower credit amounts on average. But the EIC is not designed as a straight line. It increases first and then decreases. Because of that, a weak or mixed correlation does not mean AGI is unimportant. It means the income relationship is nonlinear.

5.4 Aggregate Credit Amounts

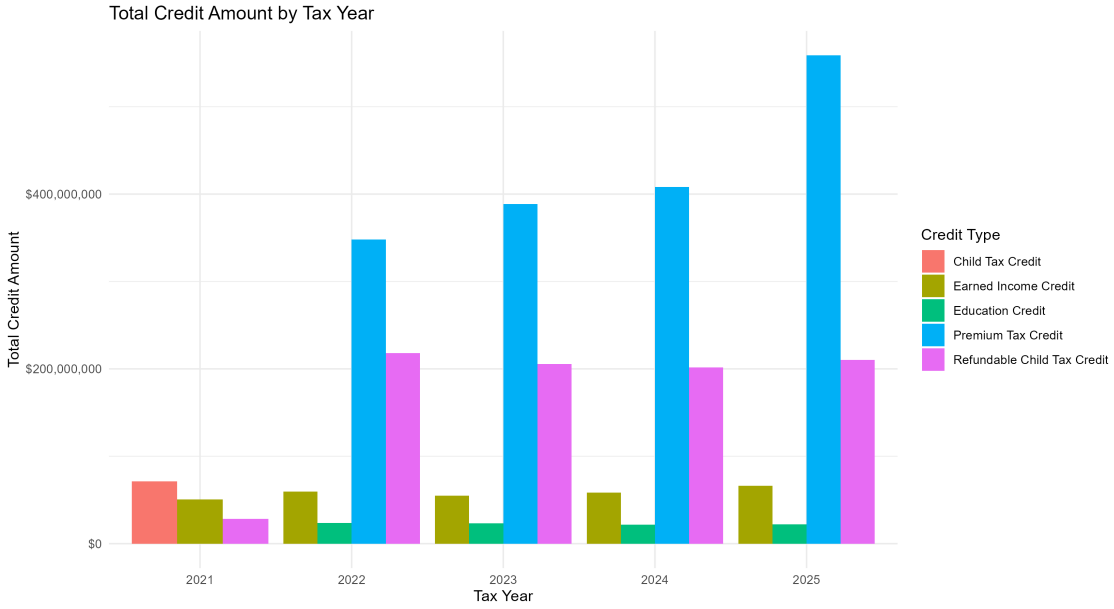


Figure 5.3: Total credit amount by tax year and credit type. Amounts are summed across AGI groups.

Figure 5.3 changes the focus from average credit per return to total credit amount. This matters because a credit can have a clear income gradient but still represent a smaller aggregate dollar amount than another credit.

The PTC is the largest aggregate credit category in the graph. It rises from roughly $350 million in the early years shown to more than $500 million in the final year shown. The RCTC is the second largest category in most years, staying near the $200 million range after the first year. The EIC is smaller in total dollars, even though it has the clearest income pattern. The Education Credit is one of the smallest aggregate categories.

This graph adds a policy interpretation. The EIC is the cleanest credit for studying

income targeting, but the PTC represents the largest total amount in this dataset. Therefore, distributional analysis should consider both average credit patterns and total federal credit dollars.

5.5 Broad Income Group Results



Figure 5.4: Average credit per return by broad income group, credit type, and tax year. Income groups are low, middle, and high based on AGI midpoint.

Figure 5.4 compares credits across broad income groups. This graph is useful because it simplifies the AGI structure into low, middle, and high income categories.

For the EIC, the low-income group generally receives the highest average credit, while the middle-income group is lower and the high-income group is mostly absent or near zero. This supports the expected EIC pattern. The EIC result is also consistent with the detailed AGI graph. The credit is not evenly spread across the income distribution. It is concentrated in lower-income groups and then declines.

For the CTC, the middle-income and high-income groups can be larger than the low-income group. This is not a contradiction. It reflects the fact that the CTC depends on qualifying children and tax liability, not just low income.

The RCTC and PTC show large values in some high-income cells. These should be interpreted carefully. Because the data are aggregated, a high average in a high-income group may reflect a small number of claimants with large credits rather than a broad benefit for all high-income taxpayers.

5.6 Correlation Analysis

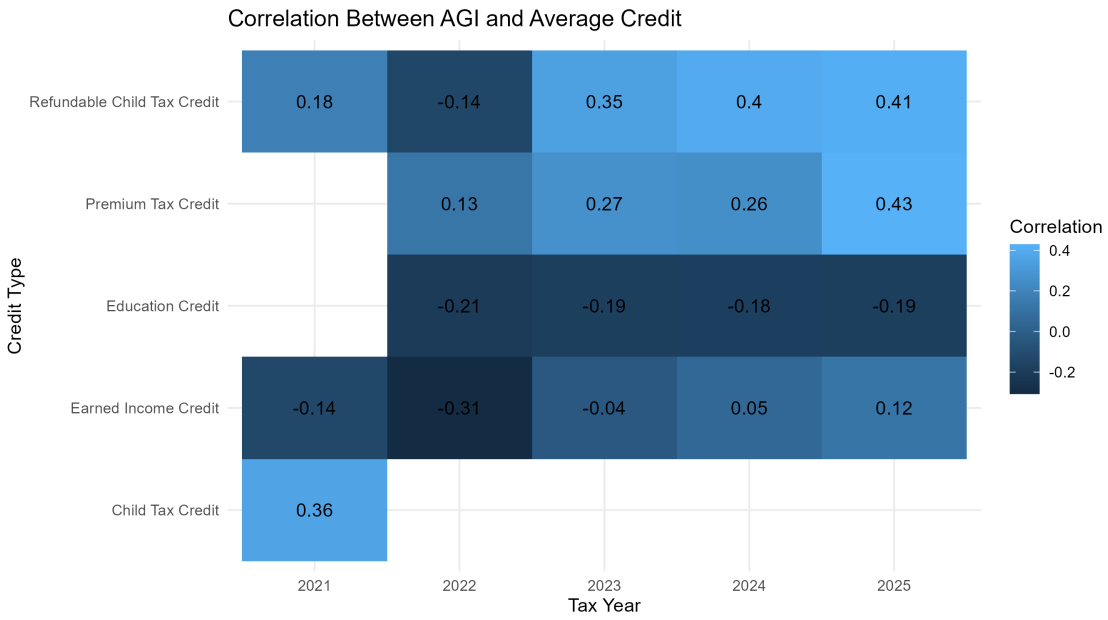


Figure 5.5: Pearson correlation between AGI midpoint and average credit per return by credit type and tax year.

Figure 5.5 reports the Pearson correlation between AGI midpoint and average credit per return. The values show whether average credit tends to rise or fall as AGI increases.

For the EIC, the correlation ranges from negative in the early years to slightly positive in later years. These values are not strongly negative in every year because the EIC is nonlinear. The credit first rises through the phase-in range and then declines through the phase-out range. Pearson correlation is a straight-line measure, so it does not fully capture that structure [14, 15].

For the Education Credit, the correlation is consistently negative in the years shown. This suggests a modest but stable inverse relationship between AGI and average education

credit. For the PTC and RCTC, the correlations are mixed or positive in several years, which shows that AGI alone does not explain every credit in a simple linear way.

Table 5.2: Pearson correlation values from the correlation analysis. Blank cells indicate no plotted value for that credit-year combination in the final figure.

Credit Type	2021	2022	2023	2024	2025
Child Tax Credit	0.36	–	–	–	–
Earned Income Credit	-0.14	-0.31	-0.04	0.05	0.12
Education Credit	–	-0.21	-0.19	-0.18	-0.19
Premium Tax Credit	–	0.13	0.27	0.26	0.43
Refundable Child Tax Credit	0.18	-0.14	0.35	0.40	0.41

The correlation results show that AGI alone does not explain every credit in a simple linear way. The EIC is the clearest example of why the graphs matter. Although the EIC correlations are not always strongly negative, the dedicated EIC graph clearly shows the expected phase-in and phase-out pattern.

5.7 Regression Analysis

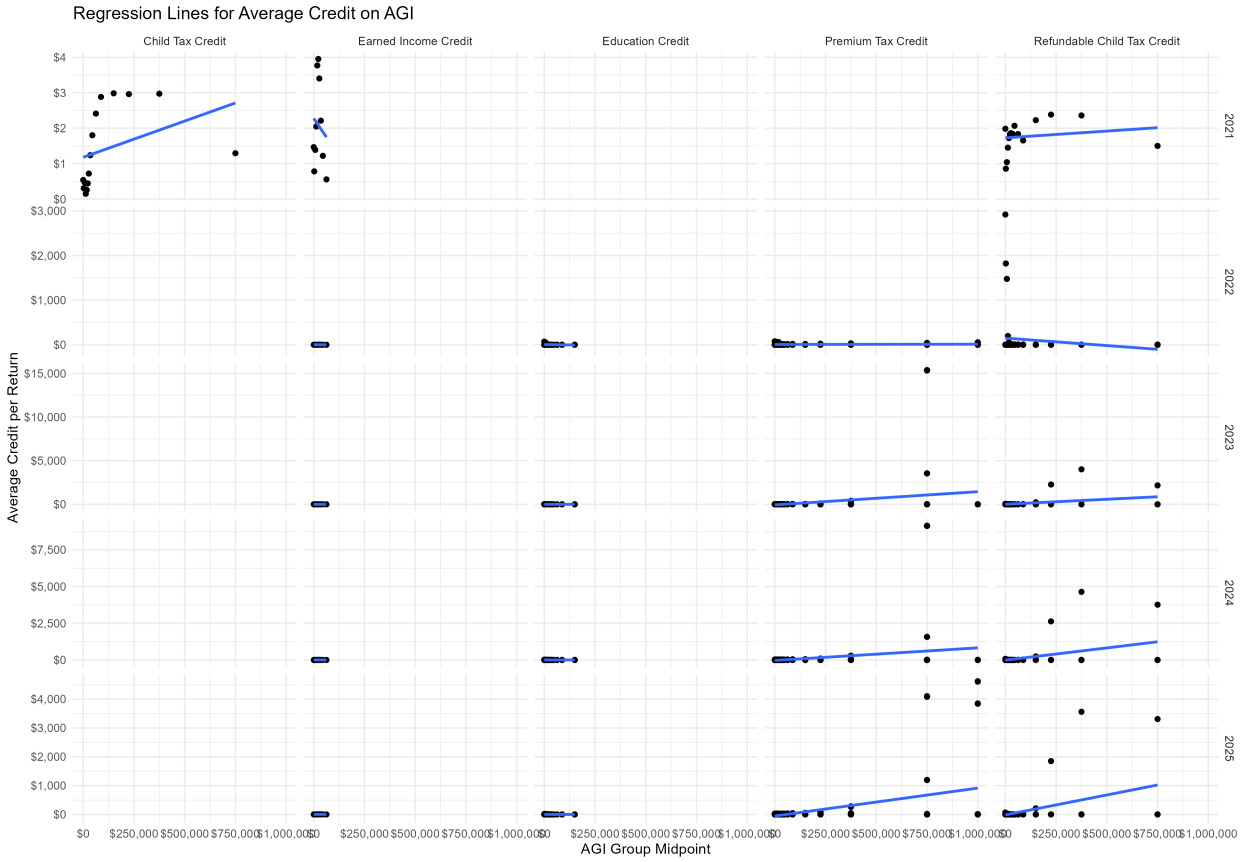


Figure 5.6: Ordinary least-squares regression lines of average credit per return on AGI midpoint, faceted by tax year and credit type.

Figure 5.6 shows simple regression lines for average credit per return on AGI midpoint. Each panel fits a line within a specific credit-year cell.

The regression results should be interpreted carefully. For the EIC, the true relationship is nonlinear. A straight-line fit does not fully capture the phase-in and phase-out structure. This is why the EIC graph is more meaningful than the slope alone. The points rise in the lower AGI groups and then fall at higher AGI groups, so a single linear slope can appear weak.

For the CTC, the regression line is positive in the year shown. This matches the graph showing that average CTC rises across lower and middle AGI groups. This is reasonable because the CTC depends on qualifying children and tax liability. It is not purely a low-income credit.

For the Education Credit, the regression lines are slightly negative in most years shown. This supports the correlation results, which showed a consistent but modest negative association between AGI and average education credit.

For the PTC and RCTC, the regression panels show substantial scatter. Some higher AGI midpoint groups have large average credits, which can produce positive slopes in some years. This does not necessarily mean higher-income taxpayers broadly receive more benefits. In aggregated data, a small number of claimants in a bracket can produce a large average.

5.8 Comparison Across Statistical Methods

The graphical analysis, correlation heatmap, regression plots, and broad income group comparison do not always give the exact same impression. This is expected because each method summarizes a different feature of the data.

The line graphs are best for identifying nonlinear patterns. They show the shape of the relationship across AGI groups. This is most important for the EIC because the credit has a phase-in and phase-out structure. The EIC line graph gives strong evidence of income targeting even when the Pearson correlation is not always strongly negative.

The correlation heatmap is best for identifying simple linear association. It is useful for comparing credit-year cells quickly. The Education Credit is the cleanest example in the heatmap because it has small but consistently negative correlations. The PTC and RCTC show positive correlations in several years, which signals that the relationship is not a simple low-to-high decline.

The regression plots are best for visualizing the fitted linear trend and the amount of scatter around that trend. They show that some credit-year cells have a smoother relationship than others. They also show when one or two high values can strongly affect the slope.

The total amount graph answers a different question. It does not describe the average claimant within an AGI group. Instead, it shows the fiscal size of each credit. This is why the PTC is important even though its average-per-return pattern is harder to interpret. It is the largest credit category in aggregate dollars.

Together, the methods support a stronger conclusion than any single method alone. The EIC is the clearest income-targeted credit. The PTC is the largest aggregate credit. The Education Credit has a modest negative AGI association. The child-related credits are more complex because family structure and refundability rules are central.

5.9 Five-Year Stability

Across the five filing seasons, the EIC pattern is the most stable visually. The level of the average credit changes by year, but the shape remains similar. The credit rises through the

lower-income groups, peaks in the $20,000 to $25,000 range, and then declines.

The total credit amount graph shows that the PTC grows substantially over time. This increase is much larger than the changes seen for the EIC and Education Credit. The RCTC remains large but relatively stable after the first year.

The correlation heatmap shows less stability because correlation only measures a straight-line relationship. For nonlinear credits, especially the EIC, the correlation is not enough by itself. Stability is clearer in the actual line graph than in the correlation coefficient.

The main five-year result is that credit behavior is persistent by credit type. The EIC remains the clearest income-targeted credit. The PTC remains the largest aggregate credit. The Education Credit remains small and modestly negative with AGI. The child-related credits remain more complicated because family structure affects eligibility and amount.

5.10 Substantive Meaning of the Results

The results have a direct interpretation for tax policy. The EIC results show that the credit is concentrated among workers in the lower-to-middle AGI range. This is consistent with the credit's purpose as wage support for low-income workers. The fact that the EIC declines after its peak confirms that higher-income groups do not receive the same average benefit.

The CTC results show that child-related policy cannot be studied using AGI alone. The credit is strongly connected to family structure. AGI matters, but the number of qualifying children and the taxpayer's liability also matter. This explains why middle-income groups can show larger average CTC values than lower-income groups.

The RCTC results show a similar issue. Since the credit is refundable, it should matter for lower-income households. However, the actual average values are uneven because eligibility requires qualifying children and because the amount can depend on earned income. This makes the RCTC more difficult to summarize with a single slope or correlation.

The Education Credit results show a small and stable negative relationship. The credit is not as large as the PTC or RCTC in total amount, and it does not show as strong a distributional pattern as the EIC. Still, its consistently negative correlations suggest that AGI is relevant.

The PTC results show why total amounts and averages must be separated. The PTC is the largest credit in aggregate dollars, but its average values by AGI group are volatile. This is likely because the credit depends on marketplace premium costs and claimant composition. A high average in one AGI bracket should not be treated as proof that the entire bracket broadly benefits.

6 Conclusions

6.1 Summary of Findings

This thesis analyzed federal tax credits across AGI groups using IRS Filing Season Statistics from filing seasons 2021 through 2025 [1]. The results show that average credit per return differs meaningfully across income groups, but the pattern depends on the credit.

The strongest finding is for the EIC. The EIC rises in the lower AGI groups, peaks around the $20,000 to $25,000 range, and then declines as AGI increases. This matches the statutory design of the credit and provides the strongest evidence that AGI-based IRS summary tables can reveal distributional credit patterns.

The PTC is the largest aggregate credit category in the data. It reaches more than $500 million in the final year shown. However, its average-per-return values are not explained by AGI alone. The credit is strongly affected by marketplace eligibility and claimant composition.

The Education Credit has a modest negative relationship with AGI. Its correlations are consistently negative in the years shown, but the average values are smaller than the major refundable credits.

The CTC and RCTC show more complex patterns. These credits depend heavily on qualifying children, earned income, tax liability, and refundability rules. Because of this, their relationships with AGI are less direct than the EIC relationship.

Overall, the thesis supports the conclusion that federal tax credits are income-sensitive, but not all credits are income-sensitive in the same way. The EIC is the cleanest example of low-income targeting. The PTC is the largest in total dollars. The CTC and RCTC require careful interpretation because AGI alone does not determine eligibility or amount.

6.2 Policy Implications

The findings show that credit design matters. Credits with explicit phase-in and phase-out rules, such as the EIC, produce clear distributional patterns across AGI groups. Credits that depend on family structure, education expenses, insurance coverage, or tax liability produce less direct patterns.

For policy analysis, this means average credit per return should not be interpreted without understanding the credit rules. A positive correlation with AGI does not automatically mean a credit is designed for high-income taxpayers. In aggregated data, a small number of claimants in a higher AGI group can create a large average value.

The EIC results are the strongest evidence that public IRS data can be used for distributional tax analysis. Even without taxpayer-level microdata, the AGI tables show the basic structure of the credit. The PTC results show the opposite side of the issue: public

AGI tables can identify fiscal importance, but they do not always fully explain eligibility patterns.

A practical implication is that analysts should use multiple summaries when studying tax credits. Average credit per return, claimant counts, total credit amount, correlation, regression slope, and broad group comparison each answer different questions. A complete interpretation needs all of them.

6.3 Limitations

There are three main limitations. First, the data are aggregated by AGI group. This prevents control for filing status, number of dependents, age, state, household size, and other taxpayer-level characteristics.

Second, AGI group midpoint is an approximation. The actual incomes within each AGI bracket are not observed. This is especially important for wide and open-ended brackets. Measurement error in the AGI midpoint can affect the estimated slope and correlation.

Third, average credit per return depends on the number of returns claiming the credit. If only a small number of taxpayers claim a credit in a specific AGI group, the average can be large and unstable. This matters most for the PTC and RCTC.

A fourth limitation is that the analysis does not separate filing status. Married taxpayers, single taxpayers, and heads of household can face different credit rules and income thresholds. Aggregating them into one AGI table can hide important differences.

A fifth limitation is that the analysis is descriptive. The results show relationships between AGI groups and average credit amounts, but they do not estimate a causal effect of income on credit receipt. A causal analysis would require individual-level data and controls for household characteristics.

Despite these limitations, the results are still useful. The EIC pattern is clear and consistent. The graphs also show where AGI alone is insufficient and where credit-specific eligibility rules must be considered.

6.4 Extensions

Future work should extend the analysis using taxpayer-level data if available. Individual-level data would allow controls for filing status, number of qualifying children, household size, and state. This would improve the interpretation of the CTC, RCTC, and PTC.

A second extension would estimate nonlinear models for credits with phase-in and phase-out rules. For the EIC, a quadratic model or piecewise regression would fit the policy structure better than a simple straight line. A piecewise model could estimate one slope for the phase-in range and another slope for the phase-out range.

A third extension would compare federal credits to state-level credits. Many states have their own earned income credits or child-related credits. Comparing federal and state credit structures would provide a broader view of how tax policy affects households across income groups.

A fourth extension would analyze credit participation separately from credit amount. This thesis focuses on average credit per claiming return. A future analysis could also study the number of returns claiming each credit by AGI group. That would separate two questions: who claims the credit, and how much do claimants receive.

A fifth extension would examine inflation and policy changes over time. The years included in this thesis include policy and economic changes that may affect credit amounts, filing behavior, and eligibility. Adjusting for inflation or separating policy years could improve the time-series interpretation.

Acknowledgments

The author thanks the Department of Mathematics and Statistics at the University of Maryland, Baltimore County for guidance throughout the Stat 497 Senior Thesis. The author especially thanks Matthias K. Gobbert for his feedback and direction during the thesis process. Experience gained through the Volunteer Income Tax Assistance program, led by Trisha Wells, provided the practical motivation for studying federal tax credits and income thresholds.

References

[1] Internal Revenue Service. Filing Season Statistics. https://www.irs.gov/statistics/filing-season-statistics.

[2] Internal Revenue Service. Statistics of Income. https://www.irs.gov/statistics/soi-tax-stats-statistics-of-income.

[3] Internal Revenue Service. Tax Statistics. https://www.irs.gov/statistics.

[4] Internal Revenue Service. SOI Tax Stats: Individual Income Tax Returns Complete Report. https://www.irs.gov/statistics/soi-tax-stats-individual-income-tax-returns-complete-report-publication-1304.

[5] Internal Revenue Service. Earned Income Tax Credit. https://www.irs.gov/credits-deductions/individuals/earned-income-tax-credit-eitc.

[6] Internal Revenue Service. Child Tax Credit. https://www.irs.gov/credits-deductions/individuals/child-tax-credit.

[7] Internal Revenue Service. Premium Tax Credit. https://www.irs.gov/credits-deductions/individuals/premium-tax-credit.

[8] Internal Revenue Service. Education Credits: Questions and Answers. https://www.irs.gov/credits-deductions/individuals/education-credits-questions-and-answers.

[9] Internal Revenue Service. Publication 17: Your Federal Income Tax. https://www.irs.gov/publications/p17.

[10] Internal Revenue Service. Publication 596: Earned Income Credit. https://www.irs.gov/publications/p596.

[11] Internal Revenue Service. Publication 970: Tax Benefits for Education. https://www.irs.gov/publications/p970.

[12] Internal Revenue Service. Publication 972: Child Tax Credit and Credit for Other Dependents. https://www.irs.gov/forms-pubs/about-publication-972.

[13] Internal Revenue Service. Form 8962: Premium Tax Credit. https://www.irs.gov/forms-pubs/about-form-8962.

[14] Devore, J. L. Probability and Statistics for Engineering and the Sciences. Cengage Learning.

[15] Diez, D. M., Barr, C. D., and Cetinkaya-Rundel, M. OpenIntro Statistics. OpenIntro.

[16] Montgomery, D. C., Peck, E. A., and Vining, G. G. Introduction to Linear Regression Analysis. Wiley.

[17] Montgomery, D. C. Design and Analysis of Experiments. Wiley.

[18] Wooldridge, J. M. Introductory Econometrics: A Modern Approach. Cengage Learning.

[19] R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. https://www.r-project.org/.

[20] Wickham, H. et al. Welcome to the tidyverse. Journal of Open Source Software, 2019.