Evaluating Head Start: The Marginal Value of Public Funds

By David Yu*

We measure the welfare impact of Head Start using the marginal value of public funds (MVPF). Head Start is a federally funded early childhood education program that provides basic services, including preschool, to children from disadvantaged families. The MVPF is the ratio of the benefits of a policy to its costs, less any externalities. Using prior studies which provide causal estimates of Head Start's effects, we compute the MVPF of Head Start and find that it has generated a large return on investment. We also find that returns are decreasing over time, consistent with the principle of diminishing marginal utility.
JEL: D61, I26, I28, I38, J13
Keywords: Head Start, marginal value of public funds

Head Start is an early childhood education program for disadvantaged students. It is an expensive federally funded program, coming in at over ten billion dollars in 2019 (Office of Head Start (2019)). There have been multiple evaluations of its effectiveness, from both government-mandated studies and quasi-experimental papers. In this article, we synthesize the results found by these evaluations, using the marginal value of public funds (MVPF) framework. The MVPF, intuitively, measures the return on investment of increasing the funding of a government policy by an additional dollar.

Using prior estimates of the causal effects of Head Start, we construct three MVPFs, corresponding to three consecutive time periods in Head Start's history, which range from inception in 1965 through 1990. We find that Head Start provides large benefits to both individuals who enrolled in Head Start and their children (regardless of whether the second generation enrolled in Head Start). Furthermore, the MVPF of Head Start is decreasing over time, from around 36 in the first five years to about 14 by 1990.

Other papers have estimated the MVPF of Head Start and of early childhood education programs more broadly. In this context, our MVPF estimates are approximately in the middle of the literature.

Section I provides background on Head Start and the MVPF. Section II reviews the literature on Head Start. Section III describes our methodology, and Section IV presents our estimates of the MVPF of Head Start. Section V discusses these results, and Section VI concludes.

* Yu: UCLA, Bunche Hall, 315 Portola Plaza, Los Angeles, CA 90095, david.yu@ucla.edu. I am very grateful to Martha Bailey for advising me on this thesis and providing extremely helpful comments.

I. Background

A. Head Start

Head Start began in January of 1964, when President Lyndon B. Johnson gave his first State of the Union address, in which he declared, "unconditional war on poverty in America." He later called for, "special school aid funds as part of our education program." Johnson's "War on Poverty" was put into motion with the Economic Opportunity Act of 1964, which contained the initial funds for Head Start. Spearheaded by the Office of Economic Opportunity (OEO), over 500,000 children participated in the first summer edition of Head Start in 1965. That fall, the first full-year programs were launched, enrolling over 20,000 students (Smith and Bissell (1970)). Head Start offered multiple forms of aid to enrolled students and their families, including preschool, health services such as vaccinations and nutrition, and mental health services.

The development of Head Start was motivated by several factors. First, the establishment and early success of model programs such as the Perry Preschool Project and the Early Training Project encouraged further development of early education programs (like Head Start). According to Smith and Bissell, "[e]ach [model program] reported that carefully designed and implemented programs increased the cognitive performance of disadvantaged children". Second, there was a growing belief that issues within the American education system could be addressed during the years just before elementary school. In particular, many studies at the time found that children of minorities and low-income families entered school already behind their peers. Third, new psychological research in the early 1960s pointed towards the malleability of intelligence, especially at early ages. Lastly, the program drew political popularity. Head Start was liked by most Americans, which helped facilitate its fast beginnings (Smith and Bissell (1970)).

Over 200,000 children enrolled in the fall of 1967. At this point, about two million students had participated in Head Start. Since then, Head Start has received repeated re-authorizations (typically for five year periods) and steady funding increases from Congress. In 1984, Head Start reached nine million total children served, and in 1994, Early Head Start was established. Early Head Start provides services similar to Head Start to low-income infants and toddlers under age 3, as well as their families. In 2000, total annual enrollment climbed to 900,000 and held steady around that level into the 21st century (About Head Start (2019)). Head Start has now served over 36 million children in total, and averages over one million enrollees per year (Administration for Children and Families (2021)).

B. MVPF

There are multiple metrics for measuring the welfare impact of Head Start. We compare the possible choices and argue that the MVPF allows for the widest use of causal estimates and facilitates easy comparisons among different government policies.

One common option is the marginal excess burden (MEB) (Hendren (2016)). The MEB associated with a tax reform is the theoretical compensation that the government should provide to an individual in order to return their utility to the pre-reform level. The MEB is positive if the tax reform lowers an individual's utility (e.g. the reform burdens them). In terms of usefulness as a welfare tool, the MEB is often used in optimal taxation analysis (Diamond and Mirrlees (1971)). But, to be empirically estimated, the MEB requires compensated effects, which can be difficult to obtain. This is because a price change has two effects on demand: the income effect (consumers adjust their behavior because their real income has changed), and the substitution effect (consumers substitute away from (toward) the good if the price goes up (down)). The compensated response isolates the substitution effect (by "compensating" for the change in real income), but this makes it hard to pin down because, in real life, the observed change is usually an uncompensated effect.

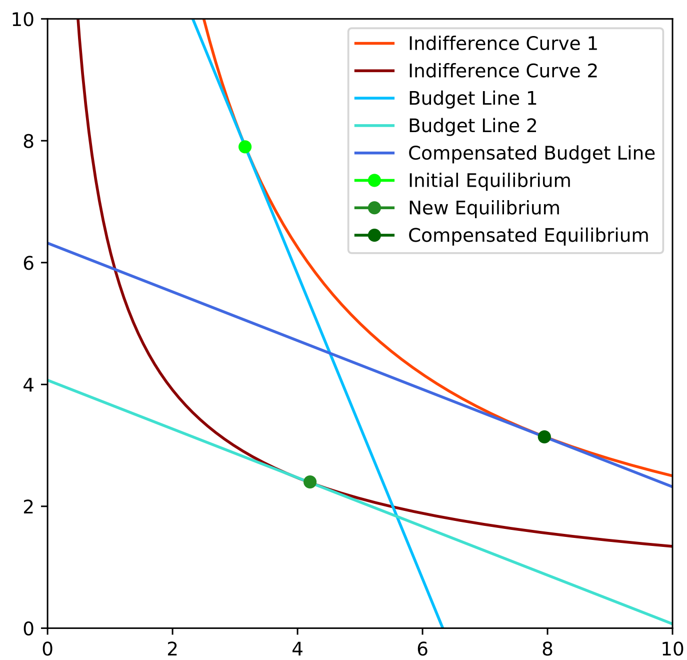

Figure 1. Compensated vs. Uncompensated Effects

Figure 1 explains this difference. Suppose that a consumer is forming their consumption bundle from two goods. The light green dot represents the initial equilibrium given the consumer's budget and utility function. After an increase in the price of the good on the y-axis, the budget line and indifference curve are each updated (from 1 to 2), and the new equilibrium point is reached. This is the effect observed in real life. The compensated effect is at the dark green point, after a theoretical return to the same real income level as before the price change. Therefore, in order to obtain compensated effects, one would need to estimate the difference between the new and compensated equilibrium points.

A further disadvantage of the MEB is that it does not account for other social costs, for example, the additional administrative cost of enforcing the tax reform. This issue is discussed by Slemrod and Yitzhaki (1996), who propose using the marginal cost of public funds (MCPF) as a way of capturing a wider set of tax change effects. The MCPF has roots in Mayshar (1990), who defines the MEB of a tax reform to be the difference of the change in consumer surplus and the change in revenue, and defines the MCPF to be the ratio of those terms (as opposed to the difference).

However, the marginal cost of public funds can be augmented by also considering the benefits to society derived from the use of the tax revenue. Slemrod and Yitzhaki (2001) advocate for the inclusion of the marginal benefit of public projects (MBP), which captures the value to society of an additional dollar spent on a particular public good. They argue for the comparison of the marginal benefit to the marginal cost, and recommend proceeding with a public project if the benefit exceeds the cost.

This concept is refined in Hendren (2016), which defines the marginal value of public funds (MVPF) of a public good as the ratio of an individual's willingness to pay for the good to the cost of the program plus any fiscal externalities. As illustrated by Hendren (2016), the MVPF does not require compensated effects, therefore making it easier to use empirically compared to the MEB.[1] Hendren (2016) also provides additional insight as to why the MCPF is lacking in comparison to the MVPF. The MCPF uses the calculated behavioral response to a theoretical modification of the tax schedule (which adjusts for the cost of raising the revenue for the new expenditure) instead of the response to the actual policy. For instance, one potential benefit of enrolling in Head Start is increased individual income (and therefore increased tax payments). This effect can be captured in the MVPF, but not in the MCPF, because of its reliance on a theoretical change in tax rates (to ensure budget neutrality). Thus, it cannot include actual changes in tax revenue (Hendren (2016)).

As a technical note, if the fiscal externalities are sufficiently large, the denominator of the MVPF may be negative. We follow Hendren and Sprung-Keyser (2020) in defining a negative MVPF to be infinity, because the policy in question is then able to pay for itself.

[1]To see an example of how to compute the MVPF, refer to Finkelstein and Hendren (2020)

The MVPF has the advantage of making comparisons between programs easy. It can be interpreted as measuring the "bang for the buck" of a government policy, in the sense that spending an additional dollar on the policy should lead to $MVPF of benefits to society. However, a particular value of the MVPF (say, 10) is not necessarily good or bad, as discussed by Finkelstein and Hendren (2020). There could be other programs with higher MVPFs, or programs with a similar MVPF but which target different segments of the population. The benefit of using the MVPF is a greater understanding of the tradeoffs associated with spending public funds, which can help inform policy decisions.

In this paper, we use the MVPF to measure Head Start's return on investment. It allows us to exploit a wider range of empirical estimates to compute it, because it relies on causal estimates, not compensated effects. The MVPF also allows for the incorporation of more of Head Start's effects on enrollees.

II. Literature

There have been numerous evaluations of Head Start, largely falling into two general categories. First, there are studies that rely on experiments, such as the Head Start Impact Study (HSIS). Second, there are quasi-experimental studies which seek out natural sources of random variation to help pin down the causal effects of Head Start. This paper collects causal estimates from these papers to construct the MVPF of Head Start.

The first study of Head Start began in 1969 as a joint effort between the Westinghouse Learning Corporation and Ohio University. The report found summer programs did not change cognitive or affective development, while full-year programs had some effect on cognitive development. However, several issues have been identified with the report's methodology. For instance, the OEO asked for studies of overall effectiveness, not more specific measures, making it difficult to draw precise conclusions. Furthermore, the selection procedure was flawed, and the study did not match Head Start and non-Head Start participants for comparison (Smith and Bissell (1970)). As part of the 1998 reauthorization of Head Start, Congress asked for another study of the effectiveness of Head Start, which resulted in the Head Start Impact Study (HSIS), a randomized control trial of about 5,000 individuals. The results of the trial found minor cognitive improvements that "faded out" by kindergarten, and no non-cognitive changes (Kline and Walters (2016)).

These early evaluations focused mainly on IQ and test scores in an attempt to measure cognitive improvement caused by Head Start. A separate strand of literature employs quasi-experimental methods to evaluate Head Start. These studies often broaden their results of interest to include factors such as educational attainment (e.g. high school graduation), annual adult income, and involvement with crime. Most papers use an individual-level data source, such as the Panel Study of Income Dynamics (PSID) or the National Longitudinal Survey of Youth (NLSY), in order to obtain data on these outcomes.

There are three common strategies for estimating the causal effects of Head Start. Some papers exploit the plausibly random rollout of Head Start. Due to the rush to get Head Start off the ground, some counties started offering Head Start before others, but not for a systematic reason. This generated random variation in exposure to Head Start among otherwise similar individuals. Other papers use the extensive information gathered in surveys like the PSID and NLSY to compare children who enrolled in Head Start with their siblings (who did not enroll). This alleviates concerns about the influence of family fixed effects on future outcomes. Finally, some studies take advantage of the OEO's extra funding effort directed towards the poorest 300 counties. This sharp cutoff created a funding discontinuity, so that if individuals who lived in counties close to the cutoff experience significantly different outcomes, then it is likely this change is due to the increase in Head Start funding Ludwig and Miller (2007).

Table 1 shows some basic information about the studies of Head Start relevant to computing the MVPF. Column 3 shows the years in which the study's participants were born. In our analysis, we will group studies with similar birth ranges, in order to obtain a more complete picture of each cohort.

Table 1—Comparison of Prior Head Start Studies

Paper	Data Source	Birth Range	Strategy
Thompson	NLSY	1957-1964	rollout
Barr, Gibbs (BG)	NLSY	1960-1964	rollout and 300
De Haan, Leuven	NLSY	1960-1964	partial ID
Johnson, Jackson (JJ)	PSID	1950-1976	county funding
Ludwig, Miller (LM)	NELS	1974*-1976*	300
Garces et al. (GTC)	PSID	1964-1977	sibling
Bailey et al. (BST)	Census/ACS	1950-1980	rollout
Deming	NLSY	1980*-1986	sibling
Currie, Thomas (CT)	NLSY	1980*-1986*	sibling
Carneiro, Ginja	NLSY	1977-1996	income RD
Kline, Walters (KW)	HSIS	2002*-2003*	experimental

Note: The birth range for BG is for the mothers surveyed in the NLSY. RD stands for regression discontinuity. A * means that the number is approximated from the paper.

We will use the estimates created by these papers to construct the MVPF. This summarizes a wealth of information into a single number with a simpler interpretation: the return on investment to society of investing an additional dollar in Head Start. Below, we review specific papers and the outcomes they study. We categorize studies by the individual-level data set they use. For an overview of Head Start studies organized by methodology, see Bailey, Sun and Timpe (2021).

The NLSY is used by Barr and Gibbs (2019), Deming (2009), and Thompson (2018). Garces, Thomas and Currie (2002) and Johnson and Jackson (2019) use the PSID. Finally, Bailey, Sun and Timpe (2021) combine data from the Census and Social Security Administration (SSA), and Kline and Walters (2016) make use of the experimental Head Start Impact Study (HSIS).

Thompson (2018) uses participants in the NLSY79 who were born between 1957 and 1964. He combines this with county-level funding data, and finds improvement in three areas of adult outcomes: education, income, and health. Specifically, exposure at preschool ages to an average-sized Head Start program yields 0.125 additional years of education and an annual increase of $2,199 (in 2012 dollars) in adult income. The advantage of Thompson's study is that it follows participants further into their lives, up to age 48.

Deming (2009) studies the children of mothers who were surveyed in the original NLSY79. These children largely enrolled in Head Start between 1984 and 1990, and participated in the Child NLSY (CNLSY). He follows these participants up to 2004, at which point they are at least 19 years old. Deming uses a sibling-based framework to account for family fixed effects, and finds significant benefits to a variety of adult outcomes, including high school graduation (8.6% increased probability) and health (7% less likely to suffer from poor health). Notably (in comparison to other studies), he does not find a significant reduction in criminal activity.

Barr and Gibbs (2019) combine the NLSY79 with data on the availability of Head Start (from the Community Action Programs and Federal Outlay System). They study both effects on the female enrollees into Head Start and their children. This differs from Deming (2009) in that there is no restriction that the children of these female enrollees also enroll in Head Start. They capture intergenerational effects by capitalizing on the plausibly random rollout of Head Start, and supplement this with a second strategy following the policy discontinuity framework. They find a positive impact on the second generation in several outcomes, such as a 13.9% increased chance of graduating high school and a 15% lower chance of criminal activity (arrest, conviction, or probation).

Garces, Thomas and Currie (2002) uses a supplement to the PSID specifically designed for this study. The 1995 PSID has a retrospective question on Head Start enrollment. They adopt a sibling comparison design, similar to Deming (2009), and study four outcomes: high school graduation, college enrollment, earnings in early twenties, and crime, finding benefits in each category depending on the race of the enrollee. For instance, white enrollees have a 20.3% greater chance of graduating high school, and African-American enrollees are 11.6% less likely to be booked or charged with a crime.

Each of these four papers investigates some components of the welfare impact of Head Start. In our analysis, we use their results on income and criminal activity as key parts of the MVPF, to provide a more complete picture of the cost-effectiveness of Head Start. Using all of these studies allows us to observe these estimated effects over a longer time frame, from the beginning of Head Start all the way to 1990.

Not all studies use the PSID and NLSY, however. Johnson and Jackson (2019) focus mainly on the synergies between access to Head Start and increased funding to public K-12 schools. However, they also study Head Start and its standalone effects. They find benefits in three areas (education, income, crime) for children who come from poor families (as an example, about a 10% increase in the probability of graduating from high school), and minimal effects for children from non-poor families.

Bailey, Sun and Timpe (2021) use a much larger sample than other papers (about 22.5 million children) by using the Census/American Community Survey in conjunction with the Social Security Administration (SSA) Numident records. They study long-run effects of Head Start and find positive effects on educational attainment (2.7% increased chance of high school graduation) and economic self-sufficiency. We will use these findings in the first step of our analysis, which measures wage gains from Head Start.

Finally, other papers have estimated the MVPF of Head Start before. Kline and Walters (2016) use the HSIS and focus on the fact that Head Start generated an externality by reducing costs in other publicly funded preschool programs (since fewer children enrolled in them). They find that considering this effect increases the positive effects of Head Start. They estimate the MVPF of Head Start to be 1.84. Hendren and Sprung-Keyser (2020) draw upon the various studies just reviewed (specifically, Johnson and Jackson (2019), Ludwig and Miller (2007), and Kline and Walters (2016)) and construct three different MVPFs, ranging from 0.72 to infinity.

III. Methodology

In our analysis, we will estimate the MVPF of Head Start during three periods: the initial four years (as Head Start was rolled out), then until about 1980, and then up to 1990. We then compute each MVPF using effect estimates from studies with samples in those time periods. If there are competing estimates within the same time range, we choose the one which relies on a sibling-based framework. The main reason is that this strategy eliminates family fixed effects. Also, unlike the 300-county method, results from the sibling strategy are more general. We now discuss how each component of the MVPF is computed.

Following Finkelstein and Hendren (2020), the MVPF is calculated as

$$\text{MVPF} = \frac{\text{Benefit}}{\text{Cost - Fiscal Externalities}}$$

where a positive fiscal externality (e.g. increased tax revenue) decreases the denominator and increases the MVPF.

The benefit is defined as a recipient's willingness to pay for the good. Since providing Head Start services is an in-kind government transfer, not cash, estimating the benefit presents some challenges. As outlined by Finkelstein and Hendren (2020) and executed by Kline and Walters (2016), one method is to estimate the effects of the treatment and then convert them to monetary values. An individual should be willing to pay for Head Start at least as much as they expect to increase their income by. We place all additional non-monetary effects in the fiscal externality component, so that recipients of Head Start are not willing to pay any more than this expected increase.

We pursue two strategies for estimating wage increases. First, if projections of income changes are available in the relevant studies, then we use those directly. Second, if those are not available, we use estimates of the change in high school graduation rates. To connect these changes in education to changes in wages, we use the following process, proposed by Deming (2009). The expected annual change in income is δ, defined by

$$\delta = e^{\mu+\epsilon} - e^{\mu}$$

where μ is the average log wage, and ϵ is the predicted log change in income. Depending on the available results, we either use an estimate of ϵ directly, or else we define it by

$$\epsilon = h \times w$$

where h is the percentage effect of attending Head Start on graduating high school, and w is the percentage effect of graduating high school on wages.

The cost is computed from official Head Start figures on annual enrollment and budget allocation. After converting all figures to 2021 dollars, we then consider a specific range of years, which is determined by the sample used in the corresponding study. We sum up the total real cost over all of those years, and then divide by the total number of enrollees in the same time frame, giving the average cost of enrolling one student in Head Start for one year.

For the final component, there is ample evidence that Head Start generates externalities in several areas. Examples include health (Ludwig and Miller (2007)), crime (Garces, Thomas and Currie (2002), Barr and Gibbs (2019)), and income taxes (Bailey, Sun and Timpe (2021)). However, numerical estimates can be hard to come by, especially in the case of health outcomes. In our analysis, we use estimates of fiscal externalities whenever it is possible to convert them to dollar values. In the appendix, we explain how we estimate the monetary benefit from changes in criminal activity.

Each externality that we do not include in the MVPF biases our estimate, because we have implicitly treated it as not being affected by Head Start. Since the estimates of changes in crime and tax revenue are positive (i.e. a reduction in crime and an increase in tax revenue), those two externalities indicate that the true MVPF will be higher than an MVPF computed without them. However, health outcomes are not as clear. Ludwig and Miller (2007) find that increasing Head Start funding results in much lower child mortality rates. These children could end up creating a positive fiscal externality (for instance, by working and paying income tax), or they could create a negative one (perhaps they survive but require government assistance to maintain their health). A similar story applies to teen pregnancy. Barr and Gibbs (2019) find that enrollees in Head Start are less likely to experience a teen pregnancy. But having fewer teen pregnancies does not clearly push the MVPF in one direction. Thus, we expect that MVPFs computed without all fiscal externalities are lower than the true MVPF (because of the crime and tax revenue) and also somewhat less precise (because of the uncertain health effects).

IV. Results

Using this methodology, we compute the benefits, costs, and externalities of Head Start, for each of our three cases. Figure 2 shows the benefits portion of each estimate, which is the expected annual increase in wages.

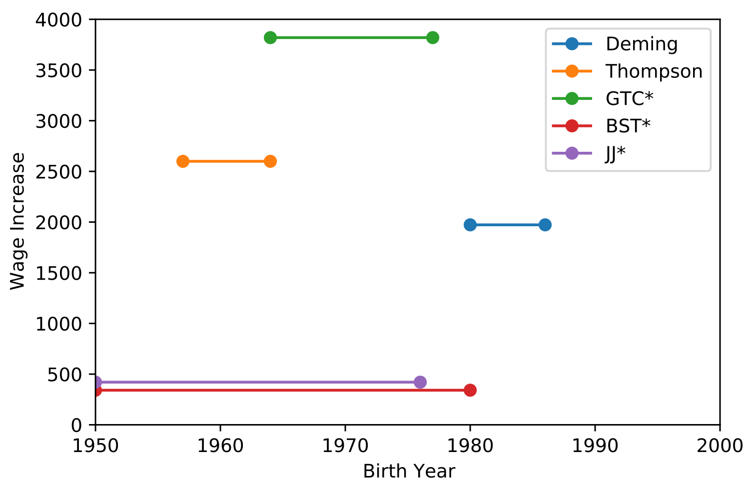

Figure 2. Annual Change in Wages (2021 dollars)

The x-axis shows the birth years of each cohort that we analyze. This figure shows that average wage gains are between $2,000 and $4,000, and that we have coverage over the first 25 years of Head Start. Note that some lines precede the introduction of Head Start in 1965. This is because those samples include individuals born just before the eligibility cutoff to make comparisons with treated individuals. We use the full line to show the total coverage of each sample, but children born in the earlier part of the interval would not have experienced these predicted income gains.

Figure 3 shows the total cost of Head Start over time. We observe that costs, adjusted for inflation, have generally risen over time. Head Start received large increases (as a proportion of prior levels) in funding in 1978 and in the early 90s.

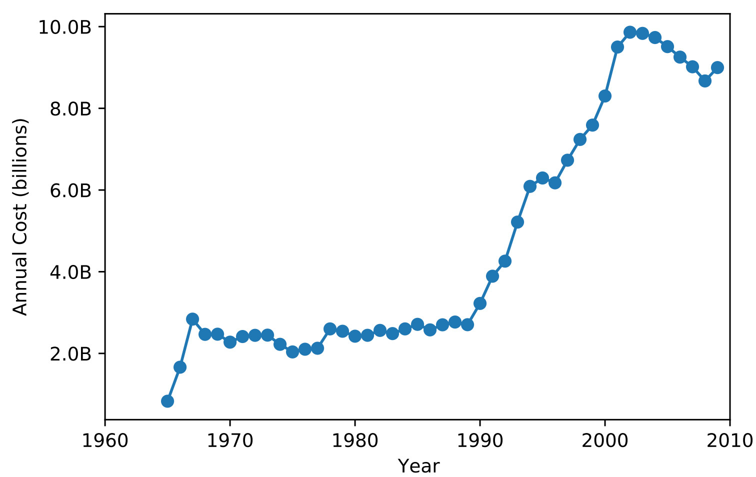

Figure 3. Total Annual Cost of Head Start (2021 dollars)

Next, we look at the average cost per student in Head Start. Total enrollment in Head Start has risen over time, as funding increases from Congress allowed Head Start to support more and more people. Figure 4 shows the average cost per enrollee. After an initial ramping up in the late 1960s, enrollment and costs increased at roughly the same pace until 1990. The early 90s saw greatly increased total funding and funding per enrollee, as Congress emphasized quality improvements in its bills to re-authorize Head Start. The Head Start Improvement Act of 1992 and the establishment of Early Head Start in 1994 are two examples.

The MVPF then ties these components together with the fiscal externalities. Our three age-sequenced estimates are given in Table 2. The years mark the approximate bounds of when each study cohort was eligible for Head Start (i.e. between the ages of 3 and 4). In the last column, we denote which fiscal externalities were considered. We do not have a criminal activity component in the first MVPF because no estimate exists for that time period, to our knowledge. Since no paper finds that Head Start increases crime rates (at any point in time), we infer that the first MVPF we compute is probably lower than the true MVPF for that time period.

From these three points, covering almost the first thirty years of Head Start's

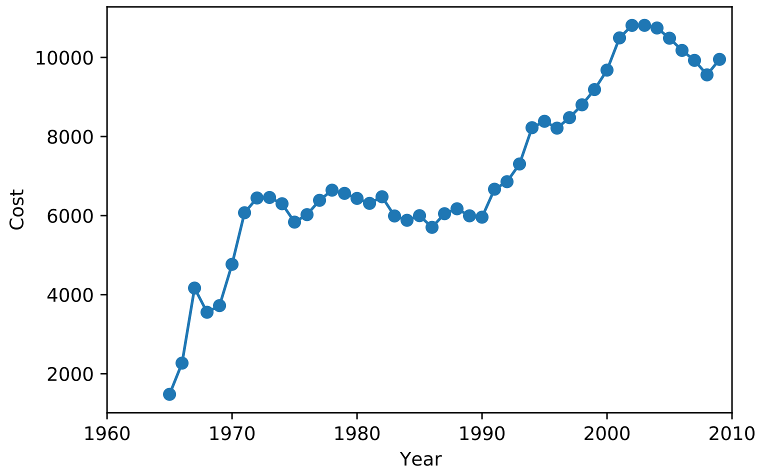

Figure 4. Average Cost per Enrollee per Year (2021 dollars)

Table 2—MVPF Estimates (by Age)

Estimate	Papers	Eligible for HS	MVPF	FEs
1	Thompson	1965-1968	35.59	None
2	GTC, LM	1967-1981	28.77	Crime
3	Deming, CT, BG	1983*-1990	14.04	Crime

existence, we observe that the MVPF is shrinking over time, though it is still relatively high in magnitude (i.e. an additional marginal investment still returns a lot). This is consistent with the law of diminishing marginal utility: as the "supply" of Head Start has increased, the marginal return is decreasing. The supply of Head Start can be interpreted to mean the number of enrollees each year, or the real annual cost of Head Start, both of which are increasing over time.

Do these estimates represent the true MVPF of Head Start for their respective cohorts? It is unlikely that they do, because the computation of the MVPF ignores many fiscal externalities. For example, Ludwig and Miller (2007) find that a "50–100 percent increase in Head Start funding reduces mortality rates from relevant causes by 33–50 percent of the control mean." It is difficult to quantify the dollar value of an additional child, and also difficult to project what their eventual impact on society will be, therefore, we don't include it in the MVPF. However, the effect of this is that our estimate of the MVPF treats health effects such as these as nonexistent. Whatever the true effect is, it will then bias our results away from the actual MVPF.

Additionally, because the first estimate does not include a crime fiscal externality, and because the general finding in the literature is that Head Start reduces criminal activity, we project a downward bias in the first estimate (e.g. it is lower than it should be). This contributes further to our observation that the MVPF is decreasing as a function of time.

V. Discussion

In this section, we compare our three results with other MVPFs computed in the literature. Figure 5 shows our results (denoted by 1, 2, and 3) against other estimates of Head Start's MVPF. In parentheses, we list the papers used to construct each MVPF, with abbreviations as denoted in Table 1. The range of each line shows the years in which the sampled individuals were eligible for Head Start. KW stands for Kline and Walters (2016). HSK stands for Hendren and Sprung-Keyser (2020). They compute three different MVPFs, using three studies. The introduction MVPF is from Johnson and Jackson (2019), the regression discontinuity is from Ludwig and Miller (2007), and the RCT is from Kline and Walters (2016) and their analysis of the Head Start Impact Study.

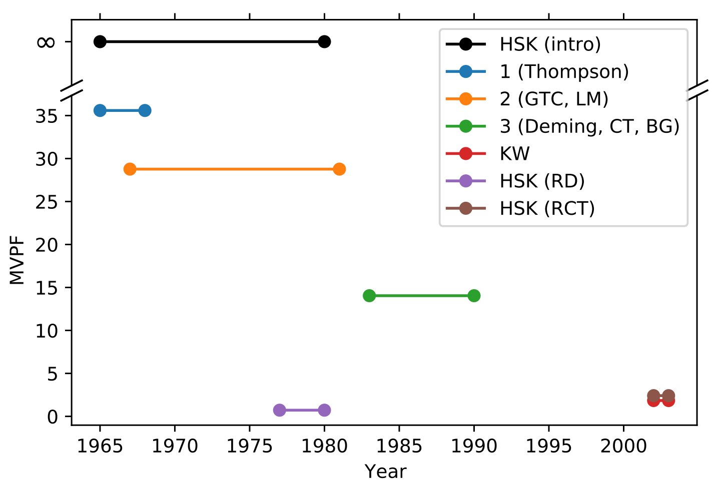

Figure 5. MVPF by Years Eligible for Head Start

It appears that the trend of decreasing MVPFs over time is continued by KW's finding and some of HSK's as well. This could be due to the rising costs of Head Start in the 1990s and into the turn of the century (in both absolute and per enrollee terms), which would increase the denominator of the MVPF. It could also simply be because there hasn't been enough time to assess the long-term benefits of Head Start to individuals who enrolled twenty years ago. If this were the case, the numerator of the MVPF for more recent cohorts might still rise as those individuals progress through their careers.

Next, we compare our preferred estimate (which is Estimate 3, based on the most recent cohort) to estimates of other early childhood education programs (from Hendren and Sprung-Keyser (2020)), such as Perry Preschool and the Carolina Abecedarian Project. The comparison is given in Figure 6. The x-axis indicates the beginning year of the study.

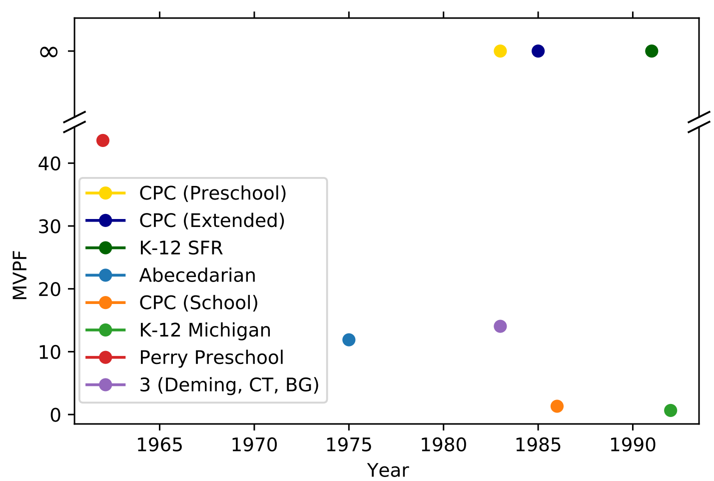

Figure 6. MVPF of Early Childhood Education Programs

As in the examination of Head Start's MVPF, we see the same downward trend of MVPF values over time. However, infinite MVPFs do exist at 1983, 1985, and 1991, so the pattern is not clear. Additionally, we observe that while Perry Preschool has an MVPF that is much higher than the MVPF of other programs, it is still close to our Estimate 1, which is valued at 35.59.

Ultimately, our estimates of Head Start's MVPF are in the same range as other estimates of Head Start and alternative forms of early childhood education. Throughout, we observe a trend of decreasing MVPFs over time, though the value is usually above 1. The trend is consistent with the principle of diminishing marginal utility, which implies that the MVPF should be decreasing because of increased funding of and enrollment in Head Start over time. The fact that Head Start's MVPF remains high is an indication that further support can still be an investment with positive returns.

VI. Conclusion

Head Start is a federally funded early childhood education program mainly targeted at children from disadvantaged families. It provides preschool and health services, and has served over 36 million children. To assess its welfare impact, we rely on estimates from the literature of Head Start's causal effects on outcomes such as adult income and criminal activity. We then unify these estimates into a single number, the marginal value of public funds (MVPF). The MVPF is our preferred welfare metric because it requires causal effects and not compensated ones. Since causal effects are easier to estimate empirically, we are able to use a wide range of empirical estimates to compute the MVPF.

We analyze the impact of Head Start through a temporal sequence that studies Head Start at three periods of time in its first 25 years. We find that the MVPF of Head Start is decreasing over time, from about 36 to about 14. Comparing these results to other estimates of Head Start's MVPF, as well as MVPFs for other early childhood education programs, we find that our values tend to fall in the middle. Some papers have found lower MVPFs (for instance, Kline and Walters (2016) obtain 1.84), and others have found higher MVPFs (Hendren and Sprung-Keyser (2020) find infinity for one of their estimates). In general, these other estimates also roughly follow the pattern of decreasing returns over time. This suggests that, as Head Start increases enrollment, the marginal benefit provided to enrolled individuals is decreasing. However, Head Start still generates a large, positive return on investment.

REFERENCES

About Head Start. 2019. "Head Start Timeline." *https://eclkc.ohs.acf.hhs.gov/about-us/article/head-start-timeline*, Accessed: 2021-11-10.

Administration for Children and Families. 2021. "Head Start History." *https://www.acf.hhs.gov/ohs/about/history-head-start*, Accessed: 2021-11-15.

Bailey, Martha J., Shuqiao Sun, and Brenden Timpe. 2021. "Prep School for Poor Kids: The Long-Run Impacts of Head Start on Human Capital and Economic Self-Sufficiency." *American Economic Review*, 111(12): 3963–4001.

Barr, Andrew, and Chloe R. Gibbs. 2019. "Breaking the Cycle? Intergenerational Effects of an Anti-Poverty Program in Early Childhood." *Journal of Political Economy (R&R)*, , (August): 1–49.

Carneiro, Pedro, and Rita Ginja. 2014. "Long-Term Impacts of Compensatory Preschool on Health and Behavior: Evidence from Head Start." *American Economic Journal: Economic Policy*, 6(4): 135–173.

Currie, Janet. 2001. "Early Childhood Education Programs." *Journal of Economic Perspectives*, 15(2): 213–238.

Currie, Janet, and Duncan Thomas. 2016. "Does Head Start Make a Difference?" *American Economic Review*, 85(3): 341–364.

De Haan, Monique, and Edwin Leuven. 2020. "Head Start and the Distribution of Long-Term Education and Labor Market Outcomes." *Journal of Labor Economics*, 38(3).

Deming, David. 2009. "Early Childhood Intervention and Life-Cycle Skill Development: Evidence From Head Start." *American Economic Journal: Applied Economics*, 1(3): 111–34.

Diamond, Peter A., and James A. Mirrlees. 1971. "Optimal Taxation and Public Production I: Production Efficiency." *Discounting and Environmental Policy*, 61(1): 8–27.

Finkelstein, Amy, and Nathaniel Hendren. 2020. "Welfare Analysis Meets Causal Inference." *Journal of Economic Perspectives*, 34(4): 146–167.

Garces, Eliana, Duncan Thomas, and Janet Currie. 2002. "Longer-Term Effects of Head Start." *American Economic Review*, 92(4): 999–1012.

Hendren, Nathaniel. 2016. "The Policy Elasticity." *Tax Policy and the Economy*, 30(1): 51–89.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *Quarterly Journal of Economics*, 135(3): 1209–1318.

Johnson, Rucker C., and C. Kirabo Jackson. 2019. "Reducing Inequality Through Dynamic Complementarity: Evidence from Head Start and Public School Spending." *American Economic Journal: Economic Policy*, 11(4): 310–349.

Kleven, Henrik Jacobsen, and Claus Thustrup Kreiner. 2006. "The Marginal Cost of Public Funds: Hours of Work Versus Labor Force Participation." *Journal of Public Economics*, 90(10-11): 1955–1973.

Kline, Patrick, and Christopher R. Walters. 2016. "Evaluating Public Programs with Close Substitutes: The Case of Head Start." *Quarterly Journal of Economics*, 131(4): 1795–1848.

Ludwig, Jens, and Douglas L. Miller. 2007. "Does Head Start Improve Children's Life Chances? Evidence from a Regression Discontinuity Design." *Quarterly Journal of Economics*, 122(1): 159–208.

Mayshar, Joram. 1990. "On Measures of Excess Burden and Their Application." *Journal of Public Economics*, 43(3): 263–289.

McCollister, Kathryn E., Michael T. French, and Hai Fang. 2010. "The Cost of Crime to Society: New Crime-Specific Estimates for Policy and Program Evaluation." *Drug and Alcohol Dependence*, 108(1-2): 98–109.

Office of Head Start. 2019. "Head Start Program Facts (Fiscal Year 2019)." *https://eclkc.ohs.acf.hhs.gov/about-us/article/head-start-program-facts-fiscal-year-2019*, Accessed: 2021-11-09.

Puma, Michael, Stephen Bell, Ronna Cook, and Camilla Heid. 2010. "Head Start Impact Study Final Report."

Slemrod, Joel, and Shlomo Yitzhaki. 1996. "The Costs of Taxation and the Marginal Efficiency Cost of Funds." *IMF Staff Papers*, 43(1): 172–198.

Slemrod, Joel, and Shlomo Yitzhaki. 2001. "Integrating Expenditure and Tax Decisions: The Marginal Cost of Funds and the Marginal Benefit of Projects." *National Tax Journal*, 54(2): 189–201.

Smith, Marshall, and Joan Bissell. 1970. "Report Analysis: The Impact of Head Start." *Harvard Educational Review*, 40(1): 51–104.

Statista. 2020. "Crime rate in the United States in 2020, by type of crime." *https://www.statista.com/statistics/202703/crime-rate-in-the-usa-by-type-of-crime/*, Accessed: 2021-12-15.

Thompson, Owen. 2018. "Head Start's Long-Run Impact." *Journal of Human Resources*, 53(4): 1100–1139.

Appendix

In this section, we estimate the cost of crime, so that we can convert a percentage estimate of the probability of committing a crime into a monetary value (i.e. converting a $x\%$ reduction in crime to an estimate of y dollars saved to society). We draw from McCollister, French and Fang (2010) for estimates of the cost to society of different types of crimes.

First, we estimate the average cost of a crime by taking the weighted sum across all crimes of the cost of that crime multiplied by its relative proportion. Second, we multiply by the chance that a random individual will commit a crime to obtain the average cost of crime per person. From Statista (2020), we take the proportions of the seven most common crimes: larceny/theft, burglary, aggravated assault, motor vehicle theft, robbery, rape, and murder/manslaughter. We then normalize each by the total number of crimes per 100,000 people (which is 2374.7). Then, we take the weighted sum as follows (the order of terms corresponds to the list of crimes above):

$$= 0.59 \times 3523 + 0.13 \times 6169 + 0.12 \times 19472 + 0.10 \times 10534 +$$

$$
\begin{aligned}
& 0.03 \times 21373 + 0.02 \times 41252 + 0.00 \times 1285146 \\
&= 2078.57 + 801.97 + 2336.64 + 1053.40 + 641.19 + 825.04 \\
&= 7736.81
\end{aligned}
$$

in 2008 dollars. Since there are 2374.7 crimes committed per 100,000 people, we have $2374.7 \times 3041 = 7221462.7$ crimes in the U.S. (3041 is the 2008 population of the U.S. in hundreds of thousands). If we make the simplifying assumption that each crime is from a separate person, this results in crime odds of 2.37% in a given year. We see that crimes are most commonly committed in the ages 20-30, so we give the average person a 23.7% chance of committing a crime ever. Thus the average crime cost per person is

$$7736.81 \times 0.237 = 1833.62$$

which is 2311.40 dollars after adjusting for inflation to 2021. For example, then, a 5% reduction in crime leads to a savings of 115.57 dollars per individual.

We recognize that there are several ways to refine this figure. We could study the changing distribution of crime over time, as the proportions of different crimes have probably not stayed the same. Depending on which crimes increase in probability, this could result in an increased or decreased estimate of average crime cost per person. Additionally, the assumption of crimes coming from separate people is unrealistic. The true cost of crime is still an active area of research and beyond the scope of this paper. We believe that this process is sufficient for our MVPF computations.