

Preliminary Draft

Who Benefits from Partnership Flexibility?

Michael Love [‡]

This Draft: October 2023

Abstract

Partnerships offer owners and investors more flexibility than any other business entity to divide up income and losses. But such flexibility also creates opportunities to reduce taxes by shifting income. There are anti-abuse rules, but these are complex and partnerships are rarely audited. The issue is important and understudied: partnerships allocated over \$17 trillion in profits and losses to partners between 2011-2020, but little is known about actual behavior. Using anonymized federal tax records from 2011-2020, I link returns to construct partnership networks, tracing profit and loss through complex entity chains, testing actual allocations against more stringent counterfactual scenarios. I estimate that flexible allocations are associated with substantial tax reduction, but that this reduction is driven by only a small set of firms that shift intensively. These firms appear to reduce taxes 7-10 cents for every dollar shifted, or roughly \$5-12 billion annually, depending on the counterfactual. They tend to be large, complex partnerships that include multiple tiers, investment funds, circular structures, profits interests, tax havens, and lots of portfolio income. The paper also sheds new light on partnership network structures and the use of carried interests.

[‡]Associate Professor, Columbia Law School, mlove@law.columbia.edu. I thank for their valuable suggestions Alan Auerbach, Emmanuel Saez, Wojciech Kopczuk, Victor Fleischer, Tom Brennan, Louis Kaplow, Steve Shavell, David Weisbach, Andrew Hayashi, Jake Mortenson, Andy Whitten, John Guyton, Matthew Smith, Gerry Auten, Lucas Goodman, Alex Raskolnikov, David Schizer, Tom Barthold, Rob Harvey, Adam Isen, Carol Wang, Nick Bull, Eric Heiser, Cecily Rock, Mark Shepard, Nirupama Rao, Rebecca Kysar, Steve Hamilton, Michael Best, Michael Smart, Jake Brooks, Jacob Bastian, Lily Faulhaber, Emily Satterthwaite, Cailin Slattery, and participants at the Harvard Law & Economics Seminar, the Columbia Tax Workshop, and the UC Irvine Tax Colloquium, and other economists at the U.S. Treasury Office of Tax Analysis. This research embodies work undertaken for the staff of the Joint Committee on Taxation, but as members of both parties and both houses of Congress comprise the Joint Committee on Taxation, this work should not be construed to represent the position of any member of the Committee.

1 Introduction

Businesses operating as partnerships represent a significant and rapidly growing share of the US economy. Partnerships are the fastest-growing class of business entity, nearly tripling in their share of business profits from only 13% to 35% since 1980 (Smith et al., 2022). They now report over \$30 trillion in assets, and over the past two decades have grown from 28% to 40% of the number of tax-reporting businesses in the United States (Black et al., 2023).

A key reason for this growth is the unique flexibility that partnerships offer owners and investors under federal tax law. In addition to being “passthrough” entities that are not subject to tax at the entity level, one hallmark feature of partnerships is the discretion they offer partners in dividing up income, gains, losses, and deductions. Unlike corporations or even other passthrough entities, partnerships make it easy for owners and investors to agree on bespoke terms that split items of income and loss very differently than their ownership shares. They may agree to allocations that are contingent, they may compensate service partners through allocations rather than using direct payments (*e.g.*, carried interest), they may modify the agreement throughout the tax year, and they may even decide to make “special allocations” to partners that deviate entirely from general patterns of the agreement.

But this flexibility is a double-edged sword: although it ideally lowers burdens on owners of a venture by making it simple to reflect the economics of their business arrangement, at the same time it incentivizes partners to engage in tax arbitrage, to make allocations that are motivated by tax-avoidance rather than the economics of the venture. Put another way, to the extent partners have different tax situations so that the same income may be taxed differently in the hands of different partners, they may be able to make allocations that exploit those differences to lower their collective tax liability. Regulations aim to prevent abusive allocations that have no “substantial economic effect” outside of tax avoidance, but there is still a range of avoidance behavior that is permissible. Also, these regulations are highly complex and involve nuanced assessments of facts and circumstances and are challenging to administer.

These observations motivate the three core questions of this paper. How much do partners utilize this allocative flexibility in practice? To what extent does this flexibility enable tax reductions? And finally, what can we say about the partnerships doing it, and who benefits?

These questions are important, for a number of reasons. First, the economic stakes are high: more than \$17 trillion of income and losses were allocated by partnerships between 2011 and 2020.¹ Second, because of their legal and economic complexity, partnerships are very rarely audited, meaning that there is little administrative oversight checking aggressive behavior (GAO, 2023). Third, partly due to regulatory complexity and partly due to lack of enforcement, com-

¹All dollar amounts are reported in 2020 dollars.

mentators have observed that “few partnership advisors are much concerned with [compliance with the rules] of allocations anymore.” (Cuff, 2009). Fourth, it is well understood that partnerships are important entities in tax planning, combining partners with highly disparate tax rates that exacerbate incentives for exploiting flexible allocation rules (Fleischer, 2008; Love, 2021). Fifth, effective tax rates for partnerships are exceptionally low compared to other entities, raising the question of how much this flexibility contributes to that outcome (Cooper et al., 2016). Sixth, and in light of these points, partnerships have been cited as likely an important component of the tax gap (Sarin and Mazur, 2023; IRS, 2023). Seventh, despite the importance of flexible allocations as a distinguishing and attractive feature of partnerships, we yet have no empirical knowledge about the utilization or the practical tax implications of this flexibility. This lack of knowledge is in no small part driven by the complexity of partnership entity structures, which make it especially challenging to trace allocations and to test counterfactuals (Cooper et al., 2016; Agarwal et al., 2021; Black et al., 2023). Finally, the distributional consequences are likely substantial, as roughly 70% of partnership income accrues to the top 1% of earners (Cooper et al., 2016).

To investigate these questions, I use anonymized federal tax records between 2011 and 2020 to estimate not only the prevalence of the use of flexible allocations, but also the apparent tax implications of such use, by comparing actual allocations against counterfactual scenarios with less flexible rules. More precisely, I take the following four steps. First, I observe actual allocations to partners reported on Schedule K-1s filed by partnerships. These information returns report the net income or loss allocated to each partner separately each year, broken down by category of income (e.g., operating income, interest, dividends, long-term capital gains, etc.). Second, I use the Schedule K-1s and Form 1065s (annual partnership returns) to link partnerships into their ownership structure chains. This is important because income and loss may pass through multiple tiers of partnerships before reaching the ultimate recipients. Third, to measure the use of flexible allocations, I compare the actual allocations to how allocations would be made under various more stringent counterfactual rules, under which tax-motivated allocations would not be possible. Finally, I link the Schedule K-1s to the annual returns of partners (e.g., individuals, corporations, trusts, etc.) to compute the tax liabilities and effective tax rates for the actual allocations versus the counterfactual.

What we are able to learn from this analysis about the use of flexible allocations depends on the counterfactual chosen, so I select three counterfactuals that each reveal distinct insights. These three counterfactuals reveal, respectively: (1) the amount of idiosyncratic “income shifting” between partners that can exploit differences in their unique tax situations, (2) the amount of profits interests or carried interests, which can be used to compensate service partners with

tax-preferred income, and (3) the extent to which any flexible allocations are used at all, compared to the more stringent rules of S-corporations, where allocations must be made pro-rata with ownership. I refer to these three counterfactuals as the *economic counterfactual*, the *no-profits counterfactual*, and the *capital share counterfactual*, each discussed in turn below.

The first counterfactual, the *economic counterfactual*, preserves the overall economic split between the partners but prevents idiosyncratic shifting of individual items of income or loss, so that partners cannot exploit differences in the character of income by advantageously allocating certain items to some partners and other items to others. To appreciate this point further, consider a hypothetical partnership with two equal partners and two types of income: a tax-exempt organization and an individual, with \$100 of taxable interest income and \$100 of long-term capital gains. Rather than splitting all items 50:50, the partnership achieves the same economics but avoids tax by allocating all of the interest to the exempt organization and all of the capital gains to the individual, who enjoys a tax-preferred rate on the latter but not the former. The economic counterfactual forces all allocations to be made consistent with the overall economic split revealed by the allocations, preventing such opportunistic picking-and-choosing. In the above example, each partner would receive \$50 each of each item. Thus, the counterfactual does not affect the overall economics of the partnership (each partner still receives \$100 in total), but simply narrowly restricts idiosyncratic allocations of particular items, and thus precludes such arbitrage.

Following the empirical methodology described above, I estimate that the vast majority of partnerships do not engage in this type of allocation shifting, but that a small number do and in fact shift quite intensively, deriving substantial tax benefits. More precisely, across all partnerships I estimate that only about 3% of all allocations are shifted away from the economic counterfactual (representing roughly \$94 billion in total shifting in 2020), reducing taxes by roughly 4% annually (about \$6 billion in 2020). But this shifting and tax reduction is not spread widely across partnerships. Instead, I find that only 3% of partnership groups shift more than 1% of their allocations in deviation from the economic counterfactual.² When measuring by dollar value, the top 1% of shifters account for 99% of the total dollars shifted. I estimate that the partnerships engaging in such shifting enjoy roughly 7-10 cents of tax reduction for every dollar shifted across 2011-2020. As discussed below, supporting evidence suggests this reduction reflects tax planning through arbitrage of the unique tax situations of partners.

Who are these high-shifting partnerships? The highest-shifting partnerships are complex, large, and have traits suggestive of tax planning. In terms of complexity, I find that arrangements involving multiple tiers and circular ownership structures are associated with especially

²By partnership group, I mean a set of related partnerships that are plausibly part of the same venture, a concept discussed in [Section 4](#).

high rates of both shifting and tax reduction. Size is also an important factor, both in terms of number of partners and income. I find that the “large partnerships” that were the subject of the recent GAO report on partnership audits (GAO, 2023) are associated with particularly high rates of both shifting and reduction. Certain traits also indicate tax planning behavior, including the presence of partners with highly disparate tax situations (e.g., having both partners with large losses and large gains from other sources outside the partnership, or highly disparate tax rates), the existence of foreign partners and partners in tax havens (which present opportunities for jurisdictional arbitrage across different types of income), the presence of lots of portfolio investment income (such as dividends, interest, and capital gains), groups involving investment funds, and the presence of “missing partners” (*i.e.*, partners that cannot be matched to annual returns).

I also investigate net flows of shifted allocations across different types of partners, different types of income, and different jurisdictions. I find strong evidence of arbitrage, with tax-preferred income on net shifted to individuals and trusts while other income is on net shifted to foreign or tax-exempt partners that is calibrated to their unique tax advantages. Taken together with the evidence above, these patterns are highly suggestive of tax planning and arbitrage.

While the no-shifting counterfactual is useful, it only tells us a narrow story of shifted idiosyncratic allocations. It does not tell us about profits interests or carried interests, which are an important aspect of partnership flexibility, or the breadth of the use of flexible allocations more generally. For this reason, I also consider the two additional counterfactuals.

The second counterfactual, the *no-profits counterfactual*, imposes the same restriction as the economic counterfactual, but *also* prevents the use of profits interests and carried interests. Profits interests are arrangements unique to partnerships, permitting partners to be compensated for services through profit allocations that would otherwise be taxable as ordinary compensation. Carried interests refer to profits interests that enable service partners to receive tax-preferred income such as long-term capital gains or qualified dividends. Thus, under the no-profits counterfactual, in addition to requiring allocations be made in accordance with the overall economic split among the partners, I also recast profits interests as compensation that generates ordinary deductions for the partnership. Note that as in the case of the economic counterfactual, the economics here are preserved—the service partners are still compensated, just with ordinary income rather than allocated profits.

Under the no-profits counterfactual, I estimate that the measured tax reduction effectively doubles: I estimate that taxes on the same reallocated income would be 7-12% higher over the actual allocations, translating to roughly \$8-12 billion annually. But even under this much broader restriction, I still find that only a narrow portion of partnerships are affected. I estimate

that although there were roughly \$120 billion flowing through profits interest and \$50 billion flowing through carried interests in 2020, only 7% of partnership groups have more than 1% of their profits flowing through profits interests, and the top 1% of partnerships account for 98% of profits interests by dollar value.

Finally, for the third and most restrictive counterfactual, I attempt to capture what can reasonably be considered as *any* use of flexible allocations unique to partnerships. This last counterfactual, which I call the *capital share counterfactual*, essentially asks: what would happen if partners were not permitted to have flexible allocations at all, and were instead required to follow rules similar to S-corporations, where allocations must be made pro-rata to the owners?³ To approximate this treatment in partnerships, I use partners’ “capital shares,” which are intended to reflect each partner’s residual or equity share of assets in the partnership, and is thus the closest analogy to the pro-rata interests of S-corporations.

Even under this most restrictive counterfactual, I find that only a narrow portion of partnerships utilize flexible allocations. I estimate that only about 20% of partnership groups shift more than 1% of their allocations. The top 1% of shifters account for roughly 92% of the shifting by dollar value. In total, only about 14% of all partnership allocations are shifted. This means that the vast majority of partnerships and partnership allocations not only adhere to the economic split of the partners, but also the capital shares. As with the first economic counterfactual, I estimate that the highest shifters again tend to be larger, more complex partnerships, with similar qualities to those described above.⁴

Four notes are important to appreciate when interpreting these results. First, the estimates are likely lower-bounds, for several reasons. The primary data source is information returns where incompleteness and errors pose challenges (see Cooper et al., 2016; Love, 2021), and where partnerships with incomplete information appear more likely to enjoy tax-reduction. Also, there is little guidance on reporting certain key fields, leaving a great deal of discretion to the taxpayer in complex situations. There is also likely attenuation of the magnitudes of the estimated tax rate changes, discussed further in Section 5. Second, it is important to recognize that these results do not (and cannot) explicitly identify illegality—although some behavior appears aggressive, the information reported is not sufficient to determine illegality.⁵

³S-corporations are a special classification of corporation for federal tax purposes that enjoy passthrough treatment, but as a trade-off are subject to strict rules to ensure especially large or public corporations cannot take advantage of the tax-privileged status. One of those rules is that they can only have one class of stock and allocations to shareholders must be made pro-rata with ownership of that stock.

⁴A tax analysis for this counterfactual is not highlighted as part of the headline results, because it is less intuitive. Under the capital share counterfactual, the economics are shifting in addition to the allocations themselves. A partner who received no income may suddenly receive substantial income, and vice versa. As a result it is less clear the economic interpretation of such analysis. I thus only discuss it briefly in the results section.

⁵Under federal tax law, allocations will be deemed improper and reallocated if they fail to have “substantial economic effect” beyond tax avoidance, but an evaluation of this legal test requires knowledge of the partner-

Third, these results should not be taken as revenue estimates—they do not measure behavioral responses to a policy change, but rather estimate (a) the use of flexible allocations under current law, (b) who appears to be benefiting most, and (c) by how much. Fourth, and perhaps most importantly of all, there are other ways to exploit partnerships that likely lead to substantially greater tax reduction opportunities than those presented by allocation shifting, notably by moving assets into and out of partnerships strategically to shift basis and defer tax. This is a topic for future research.

This paper also makes two additional contributions. First, the project provides new insights into partnership ownership structures. Although there is some nascent research on partnership structures (Agarwal et al., 2021; Black et al., 2023), there still is much to learn. As part of this investigation I complete a detailed analysis of how partnerships are interconnected. I find that partnership networks are strikingly bimodal: while most partnerships (roughly 75%) are standalone entities, and are not connected to any other partnership, a few (roughly 16%) are all interconnected into one gigantic “Mega Network” that represents the majority (roughly 60%) of all partnership income and losses. The size and inter-connectedness of this Mega Network—which consists mostly of finance firms—is naturally driven by a daisy-chaining of (even unrelated) partnership entities. Second, the paper also speaks to a longstanding and rich tax law literature concerning avoidance behavior around partnership allocation rules (Gergen, 1990, 1992; Yin, 1997; Gergen, 2003; Lokken, 2003; Fleischer, 2008; Monroe, 2010; Hasen, 2012; Monroe, 2012; Cauble and Polsky, 2014). This paper contributes to this literature by providing empirical analysis to illuminate the discussion.

2 Background: Flexible Allocations and Motivation for the Counterfactual Analysis

2.1 Flexible Allocations, Profits Interests, and Anti-Abuse Rules

Partnerships, unlike C-corporations, are not subject to entity level tax.⁶ Instead, the income, gains, losses, deductions, and credits of the partnership “pass through” to the partners, who take these amounts into account separately on their own tax returns.⁷ But while passthrough treatment is not unique to partnerships (for example, S-corporations are also passthrough business entities, and REITS, RICS, and trusts all have passthrough features), what is unique to partnerships is the flexibility in how these amounts may be allocated.

ship agreement between partners as well as the facts and circumstances surrounding the decision to make the allocations—none of which are available in my data.

⁶IRC §701.

⁷IRC §702.

Partnership allocations of income, gain, loss, deduction, or credit are determined by the “distributive shares” of the partners as specified in the partnership agreement, which need not align with liquidation interests of the partners, control rights, or other interests.⁸ The distributive shares may vary across different lines of business (if there is more than one), or across different items of income or loss. The agreement may also be amended up until the end of the tax year, adding additional flexibility, and partners may agree to make “special allocations” that deviate from partners’ distributive shares.⁹

Partnerships may also grant partners “profits interests” in the partnership, so that they receive a portion of future profits without receiving a liquidation interest in the capital of the partnership. Profits interests are generally used as a mechanism to compensate service partners, for several reasons. First, under tax law, the receipt of a profits interest is not treated as a receipt of taxable income by the service partner, because the interest is solely contingent in future profits of the partnership and does not guarantee that partner compensation. In contrast, if the partner had received a capital interests or a direct payment, that receipt would be taxable because the partner immediately has a valuable non-contingent interest in property received. Second, the profits interest creates an alignment of incentives for the service partner, very much like a commission for the success of a venture. Third, under tax law, any profits allocated to the service partner retain their character, which can generate tax savings: if an individual service partner receives long-term capital gains through the profits interest, the character of the income is preserved and the partner receives preferred tax treatment. But if the partner had instead received a direct payment for services, that would be ordinary income for the service partner.¹⁰

This enormous flexibility for partnership allocations is intentional. When the modern partnership rules were first codified in the Internal Revenue Code of 1954, a chief motivation was to preserve their unique flexibility.¹¹ For purposes of allocations, partnerships were to continue to be “mere conduits,” preserving the existing practice of allowing partners’ distributive shares and respective allocations to be determined by the partnership agreement.

Ideally, this flexibility can improve efficiency by reducing transaction costs and agency costs. In a few lines in the partnership agreement, partners can memorialize nuances of an economic arrangement that would be difficult or impossible to replicate in other entities. For instance, partnerships with multiple lines of business could align incentives by allocating a greater share of profits of a particular line to the owner-managers with greater responsibility; or partners could incentivize an investor to contribute capital by guaranteeing a preferred rate of return on

⁸Treas. Reg. §1.704-1(a).

⁹Although in such cases “tax consequences of the modification and the facts and circumstances surrounding the modification will be closely scrutinized[.]” Treas. Reg. §1.704-1(b)(4)(vi).

¹⁰See Gergen (1992) and (Fleischer, 2008) for thorough discussions of these rules and their use in practice.

¹¹S. Rep. No. 1622, 83d Cong., 2d Sess. 89 (1954). See also H.R. Rep. No. 1337, 83d Cong., 2d Sess. 65 (1954).

her investment before profits are split with other partners.¹²

Despite the potential for efficiency, this extra allocative flexibility also creates opportunities for tax arbitrage. Where partners have different tax situations, the same income or loss could have very different tax consequences in the hands of different partners. In these situations, partners have an incentive to make allocations that on net reduce tax liabilities, even if those allocations serve no business purpose. There are many dimensions over which partners have an incentive to engage in arbitrage across their tax situations to reduce their tax bill.¹³

In light of these arbitrage incentives, there are a number of complex anti-abuse rules that aim to place limits on tax avoidance, namely to ensure that partnership allocations are motivated by economics and not solely by tax reduction. In short, allocations will not be respected, and will be reallocated, if they fail to have “substantial economic effect” beyond tax avoidance.¹⁴ In practice, allocations must pass a battery of tests designed to ensure that the allocations (a) are actually reflected in the partnership agreement, (b) have “economic effect” (i.e., meaning that the allocations reflect the underlying economics of the agreement, and that there is a proper alignment between the tax and book accounting), and (c) that the economic effect is “substantial” (i.e., meaning generally that “there is a reasonable possibility that the allocation (or allocations) will affect substantially the dollar amounts to be received by the partners from the partnership, independent of tax consequences”).¹⁵ If allocations fail these tests, they will be ignored and instead reallocated to reflect a default allocation referred to as the “partners’ interest in the partnership,” or “PIP.”¹⁶ This reallocation under PIP is meant to reflect how “the partners have agreed to share the economic benefit or burden” corresponding to the tax item at stake; in other words, how the allocations would have been made if they had actually reflected the underlying economic agreement of the partners (rather than deviating from it for tax avoidance).¹⁷

It is important to clarify, however, that tax reduction alone is not *per se* abusive. An

¹²For additional discussion of transaction and agency costs with partnerships vs. corporations, see Ribstein (2009); Collins and Bey (1986).

¹³For example, partners could take advantage of disparities in their respective tax situations that include: (a) differences in the tax brackets, (b) net positive vs. net negative gross income, (c) gains or losses of different character, (d) differing amounts of deductions or net operating losses available to offset partnership income, (e) different timing of income or losses, (f) eligibility for the qualified business income deduction, or differences in the application of specified service trade or businesses under IRC §199A, (g) alternative minimum taxes, (h) net investment income taxes, (i) passive loss limitations, (j) self-employment taxes, (k) application of taxes on income to retirement accounts, (l) entity types (which may face different tax rates on the same income), (m) tax-exempt status, or (n) jurisdiction of the taxpayer (including differences in treaty rates vs. default international rates vs. domestic rates), among others.

¹⁴IRC §704(b)(2).

¹⁵See Treas. Reg. §1.704-1(b)(1)(i); Treas. Reg. §1.704-1(b)(2)(ii); and Treas. Reg. §1.704-1(b)(2)(iii). Note that there are four specific tests for substantiality, with some more strict than others. But in each case one must consider the facts and circumstances surrounding the agreement to assess the reasonable possibilities of different economic outcomes, to determine whether tax avoidance was a likely outcome. This not easy to enforce.

¹⁶Treas Reg. §1.704-1(b)(3).

¹⁷Treas Reg. §1.704-1(b)(3)(i).

allocation that lowers tax incidentally while substantially affecting the economic outcomes of the partners would not be recast under the rules described above. For example, if a partner with significant capital losses from a separate venture (and who is unable to use them) happens to be allocated a lot of capital gain while other partners happen to receive more of other types of income, it is not abusive as long as that allocation satisfies the tests laid out in the regulations as discussed above.

2.2 Counterfactual Allocations

To measure the use of flexible allocations by partnerships, there must be a standard against which to measure actual allocations. The best choice depends on the exact economic question asked. If the intent is to measure a narrower concept of shifting inconsistent with the economics of the deal, that requires a different standard than if the intent is to capture either the use of profits interests or an even broader concept of the use of flexible allocations more generally. It is for this reason that I use three standards of comparison, the *economic counterfactual*, the *no-profits counterfactual* and the *capital share counterfactual*.

Economic Counterfactual. To what extent do partnerships shift individual items of income and loss between partners inconsistent with the overall economics of the deal? To answer this question I use what I call the economic counterfactual. This counterfactual requires that partnership allocations instead be made consistent with what appears to be the economic split between the partners, so that items cannot be allocated uniquely (i.e., no slicing-and-dicing of individual items to achieve the desired overall economics). An example illustrates the concept, and how it can affect tax liabilities.

Consider two partners, a US individual and a foreign corporation. Let's say the individual is in the 37% bracket and pays 20% on long-term capital gains, while the foreign corporation is in a jurisdiction without a tax treaty, such as the Cayman Islands. Let's further say the partnership has \$100 of interest income and \$100 of dividend income to be allocated. The interest income would be taxed to the individual at 37%, but to the foreign corporation at 0% (because of the portfolio interest exemption). On the other hand, the dividends are taxed to the individual at 20% (if they are qualified dividends), but are taxed to the foreign corporation at 30% (as US income, without a treaty rate). Thus, the partners are best off by allocating all \$100 dividends to the individual and all \$100 of interest to the foreign corporation, incurring a total of \$20 tax. Under the economic counterfactual, they would need to be split proportionately to the economics of the deal (*i.e.*, a 50:50 split), so that \$50 of each item went to each partner, thus incurring \$43.50. The flexible allocations cut the rate by more than half.

The phenomenon is also possible with losses. Consider a partnership with a high-bracket

(37%) taxpayer and low-bracket (24%) taxpayer. Let's say the high bracket taxpayer has excess capital losses, and thus cannot take additional capital losses, but the low-bracket taxpayer has no capital losses. Let's say the partnership has \$100 of operating loss and \$100 of long-term capital loss. The partners do better by allocating all \$100 operating loss to the high-bracket partner and all \$100 long-term capital loss to the low-bracket partner, reducing their tax liabilities by \$61. But if the items were allocated evenly with the economic split, the total tax reduction from the losses is only \$42.50.

As these examples demonstrate, the economic counterfactual enables us to measure a concept of allocation shifting by partners relevant to the the anti-abuse regulations. The counterfactual allocation can be thought of as a very coarse approximation of PIP, discussed above; that is, a reflection of the underlying economic split. Thus, a comparison between actual allocations and the economic counterfactual reveals how much items are shifted between partners in ways that are inconsistent with the overall scheme.

But, it is important to note that although this approach captures shifting, such shifting does not necessarily fail the substantiality test. As discussed above, such a determination would ultimately depend on the facts and circumstances surrounding the agreement, and whether there was a reasonable possibility that the allocations would have a substantial effect on economic outcomes aside from tax avoidance.

No-Profits Counterfactual. As discussed above, profits interests are arrangements unique to partnerships and provide a potentially tax-advantageous way to compensate partners for their services. To what extent do partnerships take advantage of such arrangements, and what are the tax implications? To answer this question I use what I call the no-profits counterfactual. This counterfactual builds upon the economic counterfactual, requiring that allocations are made in accordance with the economic arrangement between the partners, but then also requires that profits interests be compensated as services. More specifically, the profits are allocated according to the economic interests of partners *net of* their estimated profits interests, and then the partners that would otherwise receive profits interests are compensated with a notional direct payment of ordinary income in the amount that would have been allocated under the profits interest, and the partnership is able to take an ordinary deduction that is then allocated out to the partners. Again, an example illustrates the concept, and how it can affect tax liabilities.

Consider a partnership with three partners, an individual and two tax tax-exempt organizations. Let's say the individual is in the 37% bracket and pays 20% on long-term capital gains. The individual is a service partner that is compensated with a 20% profits interest, while the exempt organizations split the remaining 80% interest. Let's say the partnership has \$200 of

capital gains. Under the actual scenario, \$40 is allocated to the service partner, generating \$8 of tax, while the remaining \$160 is allocated to the exempt organizations. Under the no-profits counterfactual, the partnership pays \$40 compensation to the service partner, generating a \$40 ordinary deduction for the partnership. The capital gains and deduction are allocated entirely to the exempt organizations, who again receive a net of \$160. The service partner is taxed \$14.80. Similar to the example above, what has changed is the character, not the economics.

Capital Share Counterfactual. To what extent do partnerships utilize the flexible allocations uniquely available to partnerships, that are not available to S-corporations? To answer this question I use what I call the capital share counterfactual. This counterfactual requires that allocations instead be made consistent with the capital shares of the partners, roughly reflecting their residual or equity interests in the assets of the partnership.

Before discussing the counterfactual further, I'll offer some background on S-corporations. S-corporations, like partnerships, are pass-through entities that are not subject to entity-level tax.¹⁸ Also similar to partnerships, S-corporation income and losses are allocated to shareholders who take these into account when calculating their own tax liabilities.¹⁹ But in contrast to the rules for partnership allocations, the rules for S-corporations are much more rigid. Allocations are made pro-rata with respect to shareholders' ownership, and owners are not able to deviate from these pro-rata shares by agreement, through profits interests, or through special allocations.²⁰ Moreover, S-corporations are permitted to have only one class of stock, and that stock must entitle the shareholders to equal rights in liquidation and distributions.²¹

Similar to partnerships, this treatment is intentional. When Congress codified the modern S-corporation rules in 1982, it was recognized that the chief purpose of S-corporations was to give business owners—notably small-business owners—an option to avail themselves of the benefits of the corporate form but without the burden of corporate tax treatment.²² The legislation thus adopted partnership-style pass-through treatment of income and losses, but without the flexibility of partnership allocations, instead maintaining pass-through on a per-share basis—a consequence of preserving the rules of a simple corporate form.²³

For this counterfactual allocation, I draw upon the closest parallel to the S-corporation pro-rata interests in the context of partnerships: the partners' "capital shares" in the partnership. Partnerships are required to maintain "capital accounts" for each partner, which record the amounts to which each partner is entitled upon the liquidation of the partnership—essentially

¹⁸IRC §1363.

¹⁹IRC §1366.

²⁰IRC §1366(a)(1) and §1377(a)(1). See also Treas. Reg. §1.1366-1(a).

²¹IRC §1361(b), and Treas. Reg. §1.1361-1(l).

²²See JCT 1982, pg. 69.

²³JCT 1982, pg. 72. Discussing how income and loss, and the character of such, would pass through as in the case of partnerships, but that the allocations would respect corporate treatment and thus would be pass through on a per-share basis.

each partner's equity share of any residual assets of the partnership. There are strict and complex accounting rules to ensure that capital accounts are maintained properly.²⁴ The capital shares are the partners' respective shares based on capital accounts, and the data is available on the Schedule K-1 for each partner.

Why compare partnership allocations to such a counterfactual? In short, by reflecting the rules applicable to the alternative passthrough form of entity, the difference will reflect the extent to which partnerships are utilizing the flexibility unique to partnerships. Again, an example can help clarify. To continue the example above for the economic counterfactual, let the nonprofit partner N have a 60% capital share and the individual partner I have a 40% capital share. Under this counterfactual, both the interest and capital gains must be split 60:40. Thus, this counterfactual is more stringent in that it both eliminates isolated shifting of individual items *and* restricts the economics to reflect the capital shares, as would be required under S-corporation rules. Thus, this counterfactual captures a broader and more general measure of the use of flexible allocations.

Limitations of the Counterfactual Approach. It's important to point out what the approach discussed above can and cannot say. These counterfactuals do allow an estimation of the use of flexible allocations, identification of which partnerships are using them, and how tax liabilities would change if the income had been reallocated in such a way. But there are important things that the analysis cannot say.

First and most importantly, the approach does not identify illegal tax evasion. As described above, tax reduction is not *per se* illegal. Allocations are illegal (and will be recast) if they fail to follow the partnership agreement or fail to satisfy the tests for substantial economic effect. In light of this fact, conclusively identifying illegal behavior is impossible on the basis of Schedule K-1 data alone. To assess illegality, one must consider the partnership agreement as well as the facts and circumstances surrounding the allocations, neither of which are available in the Schedule K-1 data.

Second, the analysis only considers a static analysis rather than a dynamic analysis with behavioral responses. Under a more restrictive allocation regime, taxpayers would undoubtedly change their behavior, and such changes would affect tax revenue. An ideal investigation would involve a natural experiment, where for some random subset of partnerships the rules permitting allocative flexibility were removed, thus permitting a comparison in the behavior between a treatment and control group. But such a natural experiment does not exist. Likewise, it is important not to consider these results as a revenue estimate. That said, these results are still deeply insightful in that they reveal which partnerships and partners are using flexible allocations and thus who would be affected by a policy change.

²⁴Treas. Reg. §1.704-1(b)(2)(iv).

Finally, it is worth pointing out that not all allocations reflect the economic scheme, but may be technical corrections to comply with safe harbors or anti-abuse provisions. Important examples include gain specially allocated with respect to minimum gain chargeback provisions or income specially allocated with respect to qualified income offsets.²⁵ These allocations will not reflect the general economic split between the partners in the year they are triggered, but will still be included in the reported allocations on Schedule K-1s.

3 Data

I use anonymized (masked) data from US federal tax records between 2011 and 2020. There are two primary sources: Schedule K-1 filings, and the annual tax returns of partners.

3.1 Schedule K-1 data

Schedule K-1s are information returns issued by partnerships that report the income, gain, loss, and deduction allocated to each partner. All US partnerships and (with only certain exceptions) all foreign partnerships with either US source income or income effectively connected with the conduct of a US trade or business must, at the conclusion of their taxable years, file both (1) a form 1065 partnership tax return, and (2) a Schedule K-1 return for every partner.²⁶ These returns are issued to the partner and filed with the IRS.

The Schedule K-1s include five key types of data that are essential for this project: (1) information on allocations to each partner, (2) capital shares, (3) profits shares, (4) de-identified (masked) Taxpayer Identification Numbers (TINs) of the partnership and the partner, and (4) dates for the beginning and end of tax years. I discuss each in turn.

Allocations. Schedule K-1s include information on the net income, gain, loss, or deduction allocated to each partner for each of the following categories: ordinary (operating) business income, rental real estate and other rental income, interest income, dividends (including ordinary, qualified, and equivalents), royalties, capital gain (short and long-term), as well as other income and deductions. Data are also available for guaranteed payments, but since these are not allocated income, I do not use them in the counterfactual analysis.

Capital Shares. Data are available on each partner's beginning and ending capital shares reported in box J of Schedule K-1, ranging from 0% to 100%. These percentages reflect the share of capital the partner would receive upon liquidation at the beginning and ending of the tax year, respectively, and are determined by the partnership agreement. Capital shares reported here are used to construct counterfactual allocations, as described in [Section 4.3](#).

²⁵See Treas. Reg. §1.704-1(b)(2)(ii)(d) and §1.704-2(f), respectively. On the other hand, section 704(c) allocations should be reported separately, and should not disturb the allocations I observe on Schedule K-1s.

²⁶See IRC §6031 and regulations thereunder, especially Treas. Reg. §1.6031(a)-1(a) and (b).

Table 1: Matching Schedule K-1s to Annual Tax Returns: Types of Returns

Domestic in **bold**, foreign in *italics*, or entities that could be either are normal font

Type of Partner	Matched Annual Return
Individual (incl. sole proprietors)	1040 , 1040-C, -F, -NR
Partnership	1065, 1065-B
Corporation	1120 , 1120-C , -H, -L, -ND, -PC, -REIT, -RIC, -S, -SF, -F, -FSC
Hybrid/Other Company	1066
Tax-Exempt	990, 990-C, -PF, -R, -T, -ZR
Trusts & Estates	1041 , 706 , 5227 , <i>3520-A</i>

Notes: This table reports the annual tax returns of individuals and entities that can be matched to Schedule K-1s, grouped by the classification of the recipient.

Profits Shares. Data are available on partners’ beginning and ending profits shares reported in box J of Schedule K-1, ranging from 0% to 100%. These percentages reflect the share of profits the partner receives according to the partnership agreement. I use the combination of profits and capital shares to estimate profits interests, as described in [Section 4.3](#).

Masked TINs. To preserve taxpayer anonymity, the K-1 returns are de-identified with encrypted, masked TINs. These masked TINs of partners reported on Schedule K-1 can be matched to other returns to determine the type of partner (e.g. individual, corporation, partnership, or other), as done by [Cooper et al. \(2016\)](#). By matching the masked TIN of the partner on schedule K-1s to the masked TIN of a different partnership—thus linking parent and subsidiary partnerships—it is possible to construct and analyze partnership networks, as discussed in [Section 4.1](#).

Tax Years. A final data element available on the Schedule K-1 is the set of beginning and ending dates for the tax period covered by the Schedule K-1. This information is crucial for understanding allocations and capital interests. First of all, partnership allocations are made at the conclusion of the partnership taxable year.²⁷ This means that one must have the correct end date or one risks matching the allocations on the Schedule K-1 to an incorrect tax year of the partner. Second, if any partner sells or liquidates their interest in the partnership in the middle of the taxable year, the taxable year ends for that partner but not the others, thus affecting the timing of the allocations and the calculation of capital accounts.²⁸

3.2 Annual Tax Returns and Entity Classifications

I follow the data methodology of [Love \(2021\)](#) to match partner Schedule K-1s to annual tax returns and to classify the entities and geographies of recipients of Schedule K-1s. [Table 1](#)

²⁷See Treas. Reg. §1.706-1(a)(1).

²⁸See Treas. Reg. §1.706-1(c).

details the types of returns that are matched to the Schedule K-1s. For Schedule K-1s that cannot be matched to an annual return, I use the inferential methodology described in [Love \(2021\)](#) to categorize the recipient by type of entity and country of tax jurisdiction.

Some entity recipients listed in [Table 1](#) are themselves passthrough entities. To estimate the effective tax rate on partnership income flowing to these entities, I also need information about the respective owners of these entities. Therefore, I also collect all Schedule K-1s issued by any S-Corporation or trust, and 1099-DIVs issued by REITs and RICs, as well as the annual tax returns of any US individuals or corporations receiving *those* information returns, and thus indirectly receiving partnership income through the intermediate passthrough entity.

The information collected from matched returns and from the inferential methodology described above is used to calculate effective tax rates on partnership income, as later described in [Section 4.4](#).

4 Empirical Strategy

To conduct the counterfactual analysis, I follow four basic steps. First, I observe the actual allocations of income and loss to partners reported on Schedule K-1s. Second, I link Form 1065 partnership returns and Schedule K-1s to reconstruct the networks of partnerships to understand how income or losses from one partnership flow up from lower to higher tier partnerships. Third, I reallocate the actual income and losses to partners following the counterfactual allocations. Fourth, I link annual tax returns of partners to estimate the taxes on the actual income versus the counterfactual income allocations.

For the no-profits and capital share counterfactuals, I must also perform an additional step as part of the reallocation. Because profits interests and carried interests are not permitted under these counterfactuals, I use the information reported on section J of Schedule K-1 to identify likely compensatory profits interests, which I instead treat as compensation to the service partners in the capital share counterfactual.

The first of these steps is straightforward, simply observing the income reported on Schedule K-1s. For the remaining steps, I describe my methodology in detail below.

4.1 Describing Partnership Ownership Structures

Under federal tax law, partners are individuals or entities that are members of the partnership, and who are entitled to certain rights to profits, control, and claims on liquidation as determined by agreement among the partners.²⁹ Partners may include all classes of tax entities: individuals, corporations, trusts, estates, and—importantly—other partnerships.

²⁹IRC §761.

Because partnerships may themselves be partners in other partnerships, partnerships can be “stacked” into complex chains through which income and loss flows. Income or loss earned by one partnership could pass through dozens or even hundreds of additional partnerships before reaching the individual or entity who ultimately reports that income on their tax return. And because there is no tax at the entity level for partnerships, this income faces no tax “leakage” as it passes between the initial partnership and the ultimate recipient, even if hundreds of partnerships lie in between.

Fortunately, because Schedule K-1s record information both about the issuing partnership and the recipient partners, the chain of partnerships (even in complex arrangements) can be identified and described using the Schedule K-1 data. More precisely, if the partner/recipient of a Schedule K-1 is itself a partnership, then the masked TIN of the recipient can be matched to a separate set of Schedule K-1s where that partnership is the issuer, and the process can be repeated. Thus, the full chain of issuers and recipients involving many partnerships can be identified and described. Even so, the complexity can be enormous as recently described and visualized in the report by [GAO \(2023\)](#).

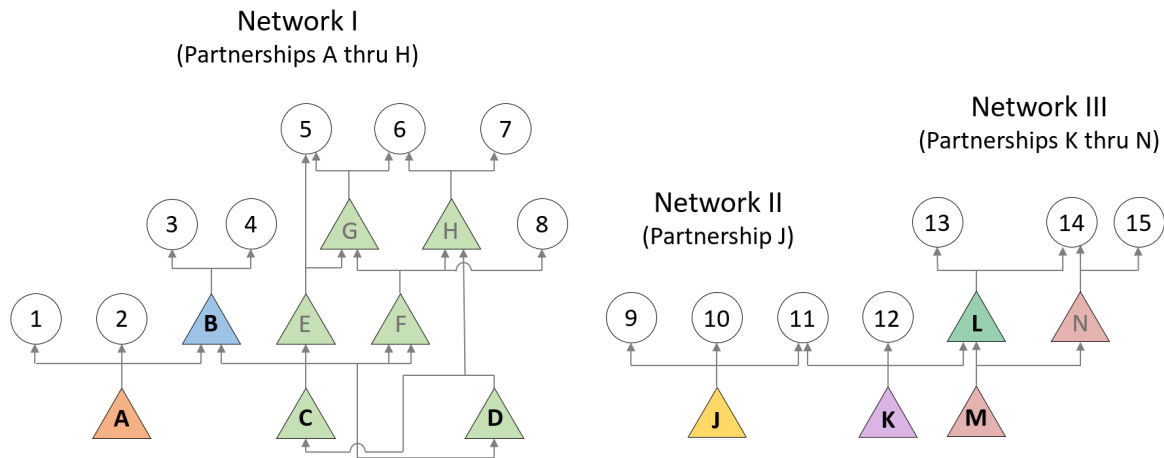
Although there are many ways to describe arrangements of linked partnerships, I choose to focus on two core concepts that I employ throughout the analysis: a partnership “network” and a partnership “group.” Notionally, I use the term network to refer to a full set of partnerships that can be connected through partner-partnership relationships. I use the term group to refer to a much narrower structure: the set of partnerships that all have a common subsidiary (or subsidiaries) among them. Below I define the terms more formally, and [Figure 1](#) provides nice examples of the two concepts.

Network: A network is a subset of the population of all partnerships, where membership in the subset includes all partnerships that have a direct partner-partnership relationship with another partnership in the subset, determined as of the end of the calendar year. ³⁰

Group: A group is a subset of partnerships in a network, where membership in the subset includes all partnerships that share the same ultimate subsidiary as all other partnerships in the subset, determined as of the end of the calendar year.

³⁰If you think of Schedule K-1s as a “bridge” that can be crossed in either direction, a network is thus the set of partnerships where any two partnerships in the set can be connected by following a unbroken path of Schedule K-1s, determined as of the end of the calendar year. Timing affects the identification and description of networks to some extent. The tax years of partners and partnerships may end for numerous reasons at times other than the end of the calendar year. As a result, the set of partners of a given partnership in June may be different than the set of partners in December. The vast majority of Networks are unaffected by this point, but it is nevertheless important to note.

Figure 1: Diagram of Example Networks and Groups:
Three Networks and Seven Groups



Notes: This diagram contains an example of three distinct partnership networks and seven distinct groups. Triangles represent partnerships. There are 13 partnerships in the diagram, labeled A thru N. Circles represent partners that are not partnerships—they may be individuals, corporations, trusts or estates. There are 15 such non-partnership partners, numbered 1 thru 15. The arrows indicate the issuance of Schedule K-1s from the partnership (the origin of the arrow) to the partner (the destination of the arrow). Partnerships may be partners in subsidiary partnerships (e.g. B is a partner of both A and C). Networks, as defined in-text, are sets of partnerships connected through Schedule K-1 issuance or receipt. Thus there are three distinct networks: (A thru H), (J), and (K thru N). Colors indicate distinct groups, defined in-text. The bold letters indicate the common subsidiary partnerships with respect to each Shifting Network, while the lighter letters indicate partnerships that are part of a Shifting Network but are not among the common subsidiary partnerships. There are seven distinct groups pictured in the diagram: (A), (B), (C thru H), (J), (K), (L), (M thru N). Thus, Network I consists of three groups, Network II consists of only one group, and Network III consists of three groups.

How are the two concepts economically useful? First, a network captures how partnerships are broadly connected through interests in one another. This is useful because it offers an indication of how economically connected partnerships are, even if they are entirely unrelated. Networks may include as few as one single partnership, or may include many partnerships. Note as well that the set of networks partition the set of all partnerships—that is, every partnership belongs to exactly one network.

Second, a group is useful because it captures the set of partnerships over which allocations can plausibly be “shifted” by partners. Because all of the partnerships in the same group share a common subsidiary, all income generated by that subsidiary could end up in the hands of any partner in the group, depending on how the allocations are agreed to. Groups are thus a much narrower concept than a network. While a network can be a sprawling web of mostly unrelated entities, all members of a group have direct interests in the common subsidiary (or subsidiaries). For this reason, I use groups as a core concept for analyzing partnership behavior. Rather than analyzing behaviors at the partnership level (which will not capture the full picture of coordinated economic activity across related partnerships), a group is a distinct

Table 2: Distribution of Group Sizes by Total Allocations, 2011-2020

Size	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
1	74.8	73.5	77.5	68.0	65.9	75.8	57.7	56.2	56.1	54.4
2	9.9	10.3	9.0	12.5	13.7	10.8	16.6	17.0	16.6	16.7
3	4.8	5.3	4.9	6.1	6.2	4.7	8.1	8.1	8.5	9.2
4	3.0	2.9	2.4	3.7	3.9	2.9	4.8	4.8	5.0	5.4
5	1.5	1.7	1.4	2.2	2.4	1.4	3.1	4.0	3.3	3.1
6	0.8	1.0	0.9	1.6	1.4	1.0	2.0	2.3	2.4	2.2
7	1.4	0.8	0.7	1.2	1.6	0.5	1.2	1.5	1.4	1.5
8	0.5	0.5	0.4	1.0	1.2	0.6	1.5	0.9	1.0	1.0
9	0.4	0.5	0.2	0.4	0.3	0.3	0.7	1.0	1.0	1.2
10-24	1.4	2.4	1.5	1.6	2.0	1.4	3.2	3.1	3.7	4.0
25-99	0.6	0.7	0.5	0.7	0.8	0.3	0.6	1.0	0.9	1.1
100+	0.9	0.5	0.6	0.9	0.8	0.3	0.4	0.0	0.0	0.2

Notes: Rows indicate the number of partnerships in the group. Point estimates in each cell represent the percentage of total allocations attributable to groups of that size.

and tractable unit that nicely reflects a set of partnerships across which partners can easily act in an economically substantive way. A group also nicely captures the “upstairs” partners and partnerships of investment funds that all have interests in a common underlying “master fund” below.

A few further observations about groups are worth mentioning. As with networks, groups partition the population of partnerships, in that every partnership belongs to exactly one group. Likewise, every group is a subset of a network, and every network can be partitioned into a set of groups. Like networks, groups may be as small as a single partnership, or may involve many partnerships. Also like networks, partners may receive allocations from more than one groups.

Table 2 shows the distribution of groups by size over time, in terms of their share of total allocations by all partnerships. The majority of allocations are made by groups consisting of only a single partnership. (Note that a single-partnership group may still be part of a larger network. But no network can have fewer partnerships than any group it contains). There has been an increasing share over time attributable to groups consisting of between 2-24 partnerships. Very little weight is attributable to very large groups. The size of the largest group varies by year, but typically falls within the low-to-mid hundreds of partnerships.³¹

A final note is that to empirically identify networks or groups in the data, I must rely on the information contained in Schedule K-1s. If a partnership is a partner in another partnership, but the issuing partnership fails to file a Schedule K-1, then it is not possible to know of the linkage between the two in actuality. Similarly, if issuer fails to accurately report the TINs of either the issuer or the recipient, the Schedule K-1 will not produce an accurate match. Finally, in the exceptional cases of foreign partnership partners not required to receive or issue Schedule K-1s, there will be no information available. In short, all network analysis thus is constrained

³¹Note, however, that groups do not necessarily overlap with the concept of a firm, and so groups will be both under- and over-inclusive when mapping tho that concept.

to the information that is actually knowable from the Schedule K-1s that are filed.

Even with these limitations, equipped with the information from Schedule K-1s and the two concepts of networks and groups described here, I find a number of new stylized facts about partnership networks. I describe them in detail below in [Section 4](#).

4.2 Identifying Profits Interests and Carried Interests

Profits interests, as described above in [Section 2](#), are interests in a partnership that entitle the holder to a portion of the profits of the partnership going forward, but that do not entitle the holder to any portion of the liquidation value of the partnership’s capital. Profits interests are typically used as a way to compensate service partners in a tax advantageous way. More precisely, service partners, rather than receiving a direct payment for their services from the partnership (which would be taxed as ordinary income), instead receive a profits interest, which entitles them to a share of the partnership’s profits. To the extent the profits take the form of long-term capital gains or qualified dividends, the partners are compensated in a tax-preferred way. This practice is well known as carried interest. Profits interests can be used in any context, but they play an especially important role in the investment industry as a way to compensate investment managers (i.e., service partners managing the capital). For a much richer and more thorough discussion of this practice, see [Fleischer \(2008\)](#).

Despite the importance of profits interests for economic and tax purposes, profits interests are not reported on any tax return. As a result, there is no way to measure the existence and use of profits interests directly. Instead, their use must be inferred from the information that is reported.

To measure profits interests I use two procedures. First, I observe the reported capital shares and profits shares reported on section J of Schedule K-1s. As a baseline, I treat any positive difference of the profit share above the capital share as a profit interests. For example, if a partner is reported as having a 20% profits share but a 5% capital share, I treat the partner as having a 15% profits interest in the partnership. Second, because the profits shares and capital shares are self-reported on Schedule K-1s, and in the context of complex payout waterfalls they may underrepresent the actual profits interest of partners, I also include alternative “backstop” measure that looks directly at allocations to certain general partners of likely investment funds. I include the excess allocations (i.e., beyond the capital shares) to the general partners of partnerships that are likely investment funds that are themselves partnerships. I determine a likely investment fund by NAICS code of the partnership, and by matching form SS-4s (applications for EINs) to the partnerships and identifying those that list being an investment fund as part of

the business purpose.³² After these amounts are aggregated, I impose a minimum size restriction to help ensure that I do not pick up spurious false positives: I exclude all profits interests that do not constitute at least 5% of the total profits share of the recipient.³³

To translate this into a tax-preferred carried interest in exchange for services, I make two further restrictions to the amount above. First, I limit the income to tax-preferred long-term capital gains and qualified dividends. Second, I restrict solely to the profits interests flowing to individuals, S-corporations, trusts, or other partnerships.

Note that this approximation is not perfect, and there are reasons to believe it is likely underestimated. First, as noted above, it is based on self-reported data on the Schedule K-1. These information returns likely receive less scrutiny than annual returns of partnerships and partners, as they do not report tax liabilities. Also, the partners may not know how to (or be able to) accurately report the distinction between profits and capital shares, especially in complex situations. The instructions are not especially detailed. Second, there is a mechanical reason these amounts are likely underestimated: capital shares will accumulate over time the extent profits are received but not distributed by the partnership. For example, if a partner starts with a 0% capital share and a 20% profits share, but then receives profits that are not actually distributed, the capital share may rise to 10% and the 20% profits interest may appear to only be a 10% profits interest.³⁴ Also as a reassurance that I am not overestimating, when looking at the industry composition in [Figure 14 in Appendix B](#), I find that essentially all of the carried interest is in the finance industry, and the little remainder is almost entirely in real estate, suggesting that I'm not capturing false-positives.

4.3 Counterfactual Allocations

To calculate the counterfactual allocations, I follow four steps. First, I eliminate duplicative income reported across multiple tiers of partnerships to prevent double-counting.³⁵ Second, I allocate this income to partners in accordance with the shares of each partners as applicable to each counterfactual. For the economic counterfactual, I allocate all items according to the partner's share of the partnership's total allocations of net profit or loss, as the case may be. For

³²More precisely, this includes partnerships that are in finance, management of companies, or real estate industries, or if the SS-4 specifies that the purpose of the partnership is investment, hedge fund, private equity fund, or similar.

³³For example, a partner with a 25% capital share and a 26% profits share would not be counted as having a 1% profits interest.

³⁴Note that there are other likely sources of imprecision, but these should not be concerning as source of systemic bias. For example, allocations may be affected by 704(c), but these could either raise or lower the estimate.

³⁵Because partnerships are pass-through entities, income passing from a lower-tier partnership to a higher-tier partnership is reported twice. To eliminate double-counting, I net out the income or losses in each partnership that were received as allocations from other partnerships. What I am thus left with is only the income either (a) generated by each partnership's activities or (b) received by each partnership from non-partnership sources.

the capital share counterfactual, I allocate all items according to each partner’s capital share reported on Schedule K-1. Third, to account for the economics of a compensatory profits interest in the case of the no-profits counterfactual and capital share counterfactual, I exclude these from the allocation shares and instead treat the partnership as making a compensatory payment to the partners in an amount equal to the profits they would have received. This payment generates ordinary income for the recipient, and an ordinary deduction for the partnership, which I in turn allocate in accordance with each partner’s applicable shares. Fourth and finally, to the extent any of the above income is allocated or paid to a higher-tier partnership, I repeat the process again, iterating as many times as it takes for all income to be allocated out to non-partnership partners.³⁶

To actually measure the utilization of flexible allocations for any given partnership, I use the following formula:

$$AmountShifted_p = \frac{1}{2} \sum_{ij} |A_{ij} - C_{ij}| \quad (1)$$

where i indexes items of income, j indexes partners in a partnership, and p denotes a particular partnership. A refers to the amount actually allocated by the partnership, while C denotes the amount that would be allocated by the partnership under the counterfactual rule. Absolute values are needed because there may be both positive and negative allocations for the same item of income in the same partnership. The calculations are divided by two because every dollar of shifted income will appear twice: first in the calculation for the partner from whom it was shifted, and second in the calculation for the partner to whom it was shifted. These amounts can be aggregated to the group level by summing across partnerships in the same group. The rate of shifting can then be calculated by taking the amount shifted divided by the sum of the absolute values of all allocations to all partners in the group.

4.4 Estimating Actual and Counterfactual Tax Rates

To estimate the effective tax rates on both the actual and counterfactual allocations, I follow the basic procedure used by Cooper et al. (2016), but with a few additional steps. As in Cooper et al. (2016), as a first step, I match the annual returns of partners to Schedule K-1s. In particular, I match forms 1040 (individuals), 1120 (corporations), 1120-F (foreign corporations), 990-T (tax exempt organizations), 1041 (trusts and estates), 1120-S (S-corporations), 1120-REIT (real estate investment trusts), and 1120-RIC (regulated investment companies). Other entities constitute a de minimis share of partnership income recipients (e.g., 1120-ND

³⁶Sometimes partnerships form “loops” or circular structures, in which case income can cycle around thousands of iterations before it is finally allocated out. This takes as many as 10,000 loops in some cases.

nuclear decommissioning funds). For entities that are themselves passthrough entities (trusts, S-corporations, REITs, and RICS) I use their information returns (Schedule K-1s or 1099-DIVs, as applicable) to match to 1040s and 1120s to identify ultimate individual and corporate beneficiaries or owners.

As a second step, I “net out” the partnership income reported on the Schedule K-1s flowing to the partner from their respective tax returns, to determine their tax situation in the absence of receiving partnership income. Likewise, for the passthrough entity partners, I net out the income pro-rata from their respective beneficiaries or owners.³⁷

As a third step, I then calculate the tax liability for each partner both with and without the partnership income. I base the calculation on available JCT calculators, adapted to the specific contexts of partnership income. I thus have a differential in tax liabilities between the scenario where the partner receives the actual partnership allocations and the scenario where the partnership income or losses are netted out.

As a fourth step, I separate out the partners that received income from partners that received deductions or losses. Receiving a deduction or loss from a partnership should generally lower one’s tax liability, while receiving income or gains from a partnership should generally increase one’s tax liability. It is important to separate out the two so that we can identify the differences between the effective rates on income versus effective rates on deductions or losses.

As a fifth step, I divide the change in tax liability by the amount of partnership income or loss received. So, for example, if we observe a partner receiving \$100 of ordinary income from a partnership, and who experiences a \$20 reduction in tax liability when that income is removed, I can calculate that income was taxed at a 20% effective rate. Likewise, if we observe a partner receive a \$100 long-term capital loss from a partnership, but observe that after removing that loss the partner only experiences a \$5 increase in tax liability, we can calculate that the partner only enjoyed a 5% effective rate of deduction on the long-term capital loss. Perhaps the partner did not have sufficient capital gains to utilize the entire capital loss.

As a sixth step, I then repeat the process above, but for the counterfactual income amounts. More precisely, after netting out the actual partnership income and losses, I then *add in* the counterfactual income and losses I calculated. Using an analogous technique, I can then estimate the effective rates of tax or deduction on the counterfactual income and losses.

A few notes are important. First, as a general rule, the goal is to trace the income until it reaches either an individual or regarded entity (e.g., C-corporation, government, tax exempt, etc.). I do not try to trace beyond regarded entities because the income does not necessarily flow beyond regarded entities. It is also a highly speculative endeavor to attempt to do so.

³⁷In the case of trusts, not all income is distributed to beneficiaries. So I assume a proportional reduction in the amounts distributed and not distributed.

Second, note that not all Schedule K-1s can be matched to returns. Ultimately, approximately 28% of income and losses flow to partners that cannot be matched to returns. This is a problem confronted by [Cooper et al. \(2016\)](#) and by [Love \(2021\)](#). Without the annual returns, the tax situations of the partners cannot be known. To estimate aggregate amounts and rates across the population of partners, I extrapolate from the partners where returns can be matched. More precisely, I first apply the methodology of [Love \(2021\)](#) to determine the type of partner (individual, corporation, trust, tax-exempt organization, government, etc.), and the jurisdiction (US or foreign). I then the average effective rates on income and loss for matched partners by type and jurisdiction (i.e., for whom I can calculate rates based on their matched annual returns), and I apply those weights to the unmatched partners. To increase specificity, I also include fixed effects for partnership size, Schedule K-1 income, industry, and network size. This procedure helps mitigate bias by using relevant available information to approximate average rates.

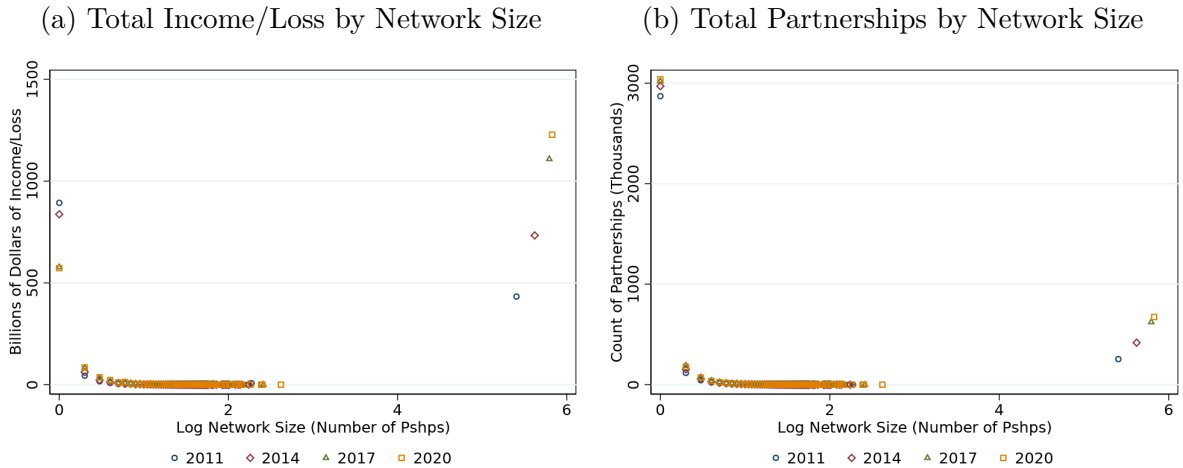
What’s the composition of the missing partners? From the inference following the procedure of [Love, 2021](#), I estimate that roughly one third of the income goes to US corporations, one third to foreign individual or entities, and roughly one third to US disregarded entities or unmatched partnerships. These missing partners are overwhelmingly in partnerships in the finance industry or that did not report an industry. The income has a similar breakdown to nonmissing income: mostly operating income and capital gains. Finally, the missing partners tend to be disproportionately from large and very large partnership groups, with total income and losses in excess of \$50 million. See [Appendix B](#) for additional details.

5 Results

I estimate that the share of allocations shifted across all partnerships is quite low, but that a relatively small number of partnerships shift intensively. I estimate that tax would be substantially higher under the economic and no-profits counterfactuals for these partnerships, and that the amount of tax reduction per dollar shifted is quite high. Even under a broader definition of flexible allocations measured by the capital share counterfactual, I still find that most partnerships do not shift allocations and take advantage of the flexible allocation rules unique to partnerships.

But before I discuss these findings in detail, I must first discuss my findings on networks, as those are prerequisites to the later estimation.

Figure 2: One Mega Network, and Many Smaller Networks



Notes: Each point represents the aggregate income (left) or the aggregate number of partnerships (right) in networks of a given size (x-axis), with different colors and shapes representing distributions in different years. For example, the first points (from the left) on each panel represent the roughly \$500bn to \$1tn (left) and about 3 million partnerships (right) that are in Networks consisting of only a single partnership. The last points on the right end of the spectrum on each panel represents the very large “Mega Network,” consisting of hundreds of billions of dollars and hundreds of thousands of partnerships. Dollar values are in 2020 dollars.

5.1 Networks and Ownership Structures

By linking partnerships into ownership chains, this investigation has yielded three new stylized facts about partnership networks that speak to questions raised by Cooper et al. (2016) and that also help expand upon findings by Agarwal et al. (2021). I discuss each in turn.

Stylized Fact 1: Partnership networks are bimodal in terms of size and structure.

On the one hand there is a single large “Mega Network” consisting of hundreds of thousands of connected partnerships and that represents the majority of all partnership income and losses. On the other hand there are millions of single-partnership networks, where each partnership is totally disconnected from all others. There are relatively few partnerships in between.

I find that partnerships are much more interconnected than previously thought. This is largely driven by a combination of (a) the popularity of partnerships as a vehicle for portfolio investment (e.g. hedge funds, private equity) and (b) the fact that partnerships are passthrough entities, so that when one partnership invests in another it quickly creates large daisy chains.

Partnerships in general are growing over time both in terms of dollar amounts and counts. In 2011, there were 3.4 million partnerships organized into 1.8 million unique networks, with a total of \$1.4 trillion of total income and losses reported to partners. In 2019, there were 4.1 million partnerships organized into 3.1 million unique networks, with a total of \$2.0 trillion of total income and losses reported to partners. This reflects the broad stylized facts that we

already know: partnerships are the fastest growing class of business entity and represent a substantial portion of total US business activity.

The patterns of networks are perhaps even more striking than the overall trends. Partnership networks tend to be bimodal, as seen in [Figure 2](#). On the one end of the spectrum, the vast majority of partnership are isolated and not connected to any other through any capital or profit interests. Roughly 3 million partnerships are in networks with only one partnership. This set represents about 30% of total income and losses of partnerships, but has been rapidly falling in share.

On the other end of the spectrum, there is the Mega Network: a network of partnerships that includes hundreds of thousands of partnerships in a single interconnected network. These are not all related—they are instead an organic development from the natural daisy-chaining of partnership interests. It is not uncommon for partnerships to hold interests in other partnerships as investments. For example, an investment fund might invest in several operating businesses that are organized as partnerships, and also invest in other investment funds, which are in turn invested in operating businesses. If the operating businesses invest in any funds themselves, then the network grows even further. Pretty quickly, the daisy-chaining creates a gigantic network.³⁸

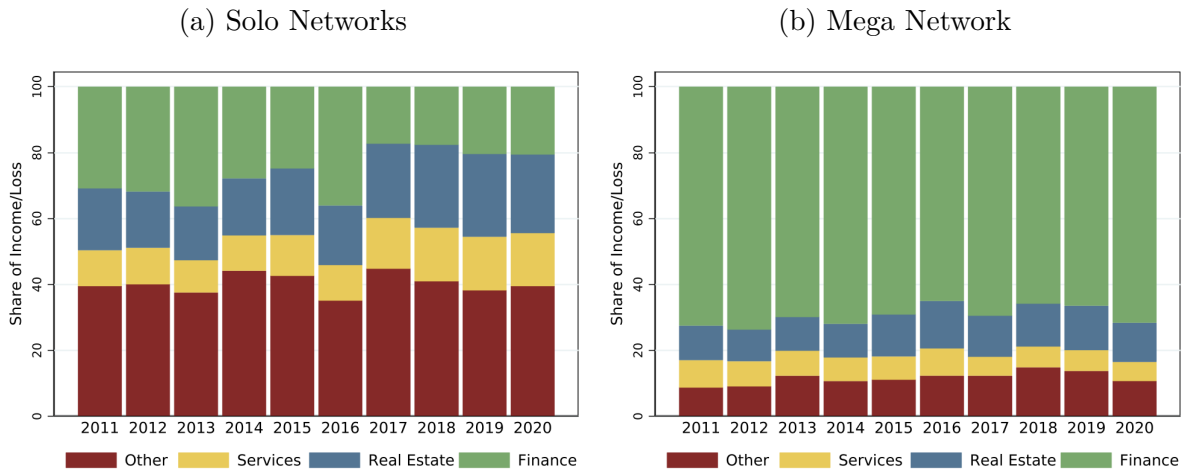
The Mega Network is also growing over time, in terms of both counts and dollars. In 2011, the Mega Network had roughly \$500 billion of income and losses and involved 250,000 partnerships. By 2020, the Mega Network had roughly \$1.2 trillion in income and losses and involved 670,000 partnerships.

Stylized Fact 2: The Mega Network is mostly finance, involves larger partnerships with more income, and income accrues in high proportions to very high earners. On the other end of the spectrum, the solo networks are smaller, tend to involve more operating businesses, and income accrues more across the distribution of earners.

The composition of the bimodal networks are different not only in terms of size and dollar amounts, but also in terms of other features like industry, income distribution, income type, and firm size. First of all, in the context of industry, the solo partnership networks are much more diverse than partnerships in the Mega Network. Roughly half of the solo networks (in terms of income and loss) are represented by operating industries that include health and professional services, or other industries like trade, agriculture, transportation, information technology, construction, or leisure and accommodation. Also, about 20% is represented by real estate—it is not uncommon for real estate rental properties to be held in standalone partnerships. Only about a quarter (and perhaps declining) of the income and loss is represented by finance. In

³⁸Note that the patterns of nodes and edges described here invite considerations in the context of graph theory, notably about what properties of the graph and of the information-generating process can be inferred from the patterns observed here, but such an undertaking is understandably beyond the scope of this paper.

Figure 3: Industry: Solo Networks vs. Mega Network



contrast, finance represents roughly 70% of the Mega Network. Operating industries represent at most about 10% of the income and loss.

The disparity between the two types of networks is also strong in terms of the types of income flowing through them. While solo networks report roughly 60% operating income or loss and 20% real estate rental income or loss, partnerships in the Mega Network report roughly 60% portfolio income or loss (including capital gain, dividends, and interest). (See [Appendix B](#)).

There is also a difference in the distribution of the income and losses across partners. For partnerships in the solo networks, roughly half of the income and losses flow to households in the bottom 99% of earners, and roughly a quarter flows to households in the bottom 90%.³⁹ In contrast, roughly 75% income and losses in the Mega Network flow to the top 1%. Roughly a third of these income and losses flow to households in the top 0.01%, compared to less than 10% of the income and losses in the solo networks.

Finally, the networks differ in terms of the sizes of partnerships as well. For solo networks, the majority of income or losses reported are by groups with less than \$5 million of income or losses. In contrast, over 60% of income or losses in the Mega Network are reported by groups with over \$50 million of income or loss, and over a third is reported by groups with over \$250 million of income or loss. (See [Appendix B](#)).

Stylized Fact 3: Discovered by [Cooper et al. \(2016\)](#), there are “circular structures” where partnerships hold interests in other partnerships that hold interests in the first partnership, creating a loop where partnerships own themselves. I find that most of these appear to be created organically by the daisy-chaining of partnership

³⁹Note that this is out of the income and losses that actually flow to households. Income flowing to other entities are not included in this analysis.

ownership.

Circular Structures have generated a great deal of interest since their discovery by Cooper et al. (2016). A question arises about what their purpose may be, when non-circular path can accomplish the same allocation of capital and returns on capital. This mystery led Guyton et al. (2021) to treat a substantial portion of the circular income reported by Cooper et al. (2016) as evasion. Can we infer something more by describing these circular structures in more detail?

First of all, what is a circular structure? I call a partnership ownership structure a “circular structure” if a partnership has an indirect profits or capital interest in itself, indirectly through its interests in other partnerships. For example, if partnership A has a 10% interest in partnership B, which has a 30% interest in partnership C, which in turn has a 5% interest in partnership A, then this represents a circular structure because A, B, and C all indirectly have interests in themselves.

I identify that the number of circular structures, the number of partnerships involved, and the amount of income and loss included in them is increasing over time. Whereas in 2011 there were roughly 350 distinct circular structures with 1,000 partnerships and about \$15 billion of income and losses, by 2020 there were roughly 700 distinct circular structures with 3,000 partnerships and roughly \$100 billion of income and losses.

The growth of circular structures is predominantly driven by the growth of a single very large interconnected network, similar to the Mega Network but just within the circular context. As seen in Figure 24, a single large circular structure accounts for essentially all of the income and loss in circular structures generally, and this has been growing rapidly. The explanation is similar to the case of the Mega Network: daisy-chaining of unrelated partnerships.

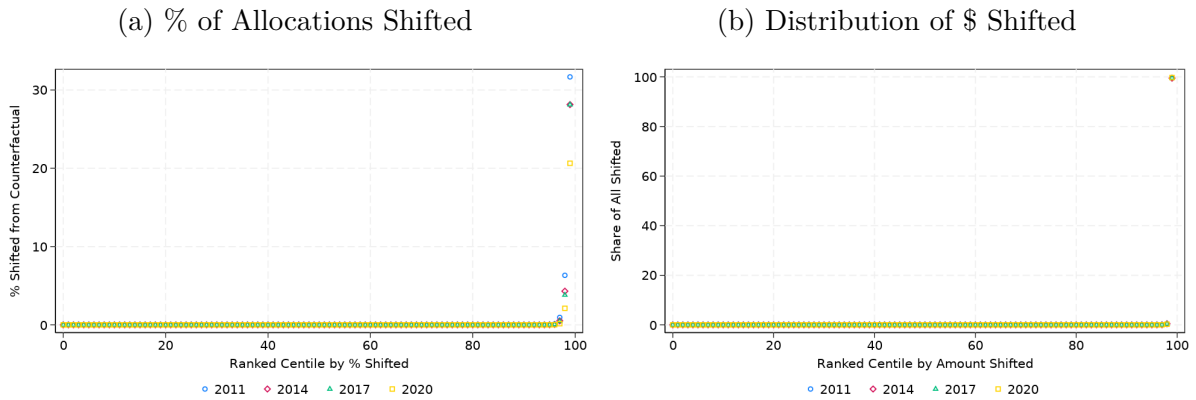
There are two important implications from this finding. First, one should not necessarily think of the circular income reported in partnerships as *per se* evasion or fraud. There is a benign explanation for the circular income: organic growth of an interconnected network of unrelated partnerships. Second, there is also the economic implication aside from the legal implication: the growth of circularity over time reflects the growing interconnectedness of partnership interests.

5.2 The Economic Counterfactual

When comparing actual allocations to the economic counterfactual, I find that only a small number of partnerships are engaged in such shifting, but that these partnerships shift intensively. They also enjoy substantial tax benefits associated with this shifting.

Shifting Behavior. I estimate that only about 3% of all partnership allocations are shifted from the economic counterfactual (roughly \$94 billion of total shifting in 2020). But this behavior is highly concentrated. I estimate that only 2-3% of partnership groups shift more

Figure 4: Shifting Under the Economic Counterfactual is Highly Concentrated



Notes: The left panel identifies partnership groups (as defined in Section 4), and ranks them (aggregated into centiles) by what percentage of their total allocations are shifted. The right panel shows the distribution of the total amounts shifted (e.g., the point at 99% means that the top centile of groups accounts for 99% of all shifted allocations)

than 1% of their allocations, but for those that do, they average shifting about 19% of their allocations. When measuring by dollar value, the top 1% of shifters account for 99% of the total dollars shifted, as the firms that shift tend to be large.

These heavy-shifting groups (*i.e.*, groups shifting more than 1% of their allocations) tend to differ from the rest of the population. They tend to be much larger in terms of income and loss, averaging \$11 million per group versus \$670,000 for non-shifters. They are also larger in terms of average number of partners (140 vs. 5). They are much more likely to be a part of the Mega Network (48% vs. 11%), and more likely to be in the finance industry (41% vs. 10%). A greater portion of their income and loss is tax-preferred as interest, dividends, or long-term capital gains (39% versus 15%). More of their partners are foreign (17% vs. 3%), are in tax havens (10% vs. 1%), are corporations (26% vs. 6%), are trusts (26% vs 6%), and are entities generally (59% vs. 21%). They have more income flowing through profits interests on average (\$640k vs. \$24k). They are also more likely to have partners receiving Schedule K-1s that cannot be matched to annual returns.

I estimate that the most intense shifting tends to occur in portfolio income. In particular, I find that dividends and interest are shifted more than any other type of income (roughly 4-5% shifted) followed by capital gains (roughly 3%). Ordinary operating income and rental income are shifted the least, at about 2%. There is also substantially more shifting at the extremes of the income distribution: the top 0.01% of earners and earners near the bottom tend to shift 3-4% of their income, while earners in the 50th thru 99th percentiles shift less than 2%. (Note that the high shifting near the bottom is attributable to partners with substantial losses, as opposed to merely low-income partners). I also find that shifting is especially prominent in the finance, manufacturing, and oil & gas industries.

I estimate a great deal of persistence of shifting behavior over time, as seen in [Table 8](#) and [Table 9](#). I estimate that roughly 99% of firms that aren't shifting relative to the economic counterfactual in a given year will not shift in the subsequent three years. By comparison, at least a third of those shifting in the current year will do so again in the next three years.


What traits are associated with the highest rates of shifting? To answer this question I run a sequence of regressions where the dependent variable is the share of group allocations shifted and the independent variable is one of a list of dummies indicating the presence of a trait. I use 60 dummy variables reflecting a broad range of qualities, including the network and ownership structures, size, type of income, types of partners, geography of partners, industry, variation in allocations, variation in capital shares, profits interests, tax situations of partners, income distribution, and lag terms for prior behavior. The dummies and their descriptions are listed in [Table 3](#), with means by count and dollar-weighted reported in [Table 7](#). I include year fixed effects to avoid time-varying shocks correlated with traits. The results with the highest estimated shifting for the economic counterfactual are presented in [Figure 5](#), where point estimates can be interpreted as the extra percentage points of total group income shifted by groups with that trait over groups without that trait.

I find that network structure, size, and opportunities for tax arbitrage are traits associated with the most prominent shifting. The trait associated with the highest shifting is large amounts being allocated to both partners with net gain and net loss tax positions, a situation ripe for arbitrage.⁴⁰ Similarly, high variation in the tax situations of partners also ranks high. But the most prominent traits overall appear to be network-related: circular structures, high tiers, groups with more partnerships, and being part of the Mega Network are all associated with high rates of shifting. Larger partnerships are also associated with high shifting rates, partnerships with 100 or more partners and over \$500 million of assets rank prominently. The “large partnership” designation that was the focus of the recent GAO audit ([GAO, 2023](#)) also has high rates of shifting. The prominence of the positive lag term reiterates the persistence of shifting across time. On the opposite end of the spectrum, the group traits associated with the least shifting are those that are the distinct opposite of the above: low tax variation between partners, small size, and simple structures, as seen in [Appendix B](#) in [Figure 27](#).

Where, on net, is this income being shifted from and to? I trace net flows between types of partners across different types of income, and I find that net flows move in tax-advantageous directions. Results are presented in [Figure 26](#) in [Appendix B](#). I estimate that over the ten years observed, roughly \$20 billion of ordinary income was shifted from individuals, grantor trusts, US corporations, tax-exempts, and foreign partners to partners that could not be matched

⁴⁰Two simple examples are timing and preferential rate arbitrage. Allocate more gains to partners in years they have losses. Or allocate tax-preferred income to gain partners and less preferred income to loss partners.

Figure 5: High Shifting Under Economic Counterfactual, by Group Trait




Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the share shifted at the group level, following Equation 1. Observations are at the group level. Point estimates represent p.p. increases in the shifting rate associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with most shifting.

to Schedule K-1s, predominantly unidentified corporations and other passthroughs or hybrid entities (like LLCs). On the flip side, I find roughly \$12 billion of long term capital gains shifted to individuals and grantor trusts, and \$3 billion in dividends shifted to tax exempt entities (avoiding unrelated business income tax). I also find about \$7 billion in interest shifted away from individuals, US corporations, and tax-exempts to predominantly foreign partners, who are eligible for the portfolio interest exemption. Although each of these shifts are tax-advantageous, the total does not appear sufficient to account for the overall tax reduction. I estimate that only about \$50-\$60 billion was shifted between types of partners in this way. This implies that most of the tax reduction from shifting comes from shifting across other more nuanced dimensions, such as rate differentials, ability to take deductions or losses, or passive loss limitations.

Rate Reduction. I estimate that tax rates on the same income are substantially higher under the economic counterfactual compared to actual allocations. I find that overall, taxes would be about \$6 billion higher annually, or roughly 4%. But once again, these benefits are highly concentrated. I estimate that only 2-3% of partnership groups enjoy a net rate reduction of at least 1 p.p., and again I estimate that the top 1% of groups enjoy over 99% of the aggregate benefit. I further estimate that these firms achieve an effective 7-10 cent reduction in tax per dollar shifted.

Which traits are associated with the highest rate reductions? As I did to investigate shifting, I run regressions with the same dummy predictors to describe the partnerships with the greatest

Figure 6: Rate Reduction Under Economic Counterfactual, by Group Trait (High)




Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the percentage point rate reduction at the group level. Observations are at the group level. Point estimates represent p.p. reduction in the effective tax rate associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with highest net reductions.

overall rate reductions. In this case, my dependent variable is the estimated percentage point rate reduction at the group level. Similar to the results for shifting, the features that stand out the most relate to the partnership’s network structure, discrepancies in the outside situations of the partners, and size. Notably, partnerships in circular structures have very high rate reduction. Complexity and size of the network is important too, as higher-tiered partnerships and those in larger groups have high rate reductions. In terms of size, partnerships with substantial income and loss and once again “large partnerships” as defined by the GAO (GAO, 2023) are associated with higher rate reductions. With regard to the partners themselves, I find that there is higher rate reduction when partnerships make substantial allocations both to partners with outside gains and to partners with outside losses, offering opportunities for arbitrage. In addition to these descriptive features, I also find that partnerships that shift more tend to have higher rate reductions, revealing a nonlinear relationship between shifting and rate reduction. I also find that rate reductions are persistent, in that enjoying prior rate reductions are associated with high rate reductions in current years.

I further investigate the traits associated with a higher intensity of rate reduction. Put another way, which traits are associated with higher tax reductions per dollar shifted? To measure this, I run the same regression above, but with a dependent variable measuring the cents of tax reduction per dollar shifted. I estimate that the highest average rates of intensity are associated with partners that indicate either tax planning or less than full transparency.

Figure 7: Cents of Tax Reduction per Dollar Shifted, Economic Counterfactual (High)



Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the cents of tax reduction per dollar shifted at the group level. Observations are at the group level. Point estimates represent p.p. reduction associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with highest net reductions.


These include groups with partners that can't be matched to annual returns, groups with large portions of income going abroad or to tax havens, groups with trust partners, groups that have an investment fund themselves or that are above an investment fund in their entity structure, and groups whose income is predominantly portfolio investment income. These high intensity partnerships also tend to be larger, and have high variations in the tax situations between partners.

On the opposite end of the spectrum I also investigate which groups have the lowest rates of net reduction, and the lowest intensity of reduction. These tend to be the opposite of the high-reduction groups, in terms of size, structure, complexity, and type of partners. They include small operating or rental businesses, partnerships with low variation in their allocations, partnership with all individual partners, partnerships with only two partners, standalone partnerships not in the Mega Network, and partnerships where the partners have similar tax situations.

5.3 Profits Interests, Carried Interests, and the No-Profits Counterfactual

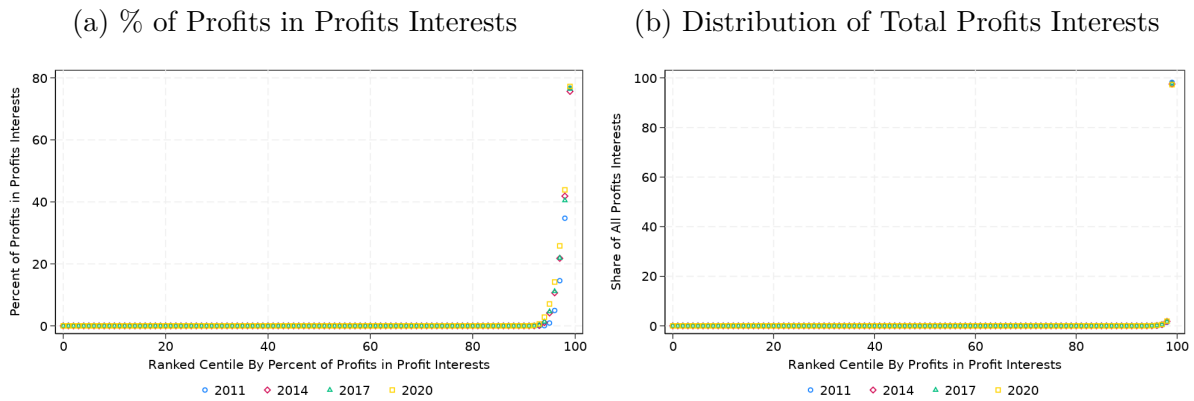
In this section I present the results for the no-profits counterfactual, but to do so must first discuss my findings regarding profits interests and carried interests. Following the procedure identified in [Section 4.2](#), I estimate that roughly 5% of total profits reported by partnerships flow to partners through profits interests, a rate that is growing gradually, as seen in [Figure 8](#).

Figure 8: Estimated Profits Interest and Carried Interest Over Time



Notes: The left panel shows the share of total partnership profits estimated to be flowing through profits interests and carried interests, over time. The right panel shows estimated amounts over time.

Figure 9: Only a Small Portion of Partnerships Utilize Profits Interests



Notes: This figure identifies partnership “groups” (i.e., distinct networks of partnerships and partners that share a common underlying subsidiary), and ranks them by their use of profits interests. The left panel shows the share of total profits in the group flowing through profits interests. The right panel shows the distribution of the total amount of profits flowing through profits interests (the point at 98% means that the top centile of groups accounts for 98% of all profits in profits interests).

Likewise, I estimate roughly 2% of total profits are flowing through carried interests, a share that is also growing over time. The total amounts are increasing as well, with now about \$120 billion appearing to flow through profits interests, while about \$50 billion flows through carried interests. Extending back to the earliest data in 2011, the past figures are on par with estimates by (Knoll, 2008), but evidence suggests substantial growth since that time. But again, for the reasons given above, this is likely a conservative estimate.

A few points are worth discussing further. First, there appears to be a substantial use of profits interests that do not generate carried interest. From figures in Appendix B, evidence suggests these are predominantly operating income, and likely split between finance, real estate, and health and services sectors. This suggests that these flows are likely in most cases either (a) profits interests related to investment, but where the investments are in passthrough entities that themselves generate operating income, or (b) profits interests in the service industry where

certain service partners are compensated with profits interests rather than becoming capital partners or becoming contractors or employees.

Second, only a very narrow class of partnerships and partners appear to utilize partnership interests. As seen above in [Figure 9](#), only about 7% of partnership groups (defined below) have more than 1% of their profits flowing through profits interests. The vast majority of partnership groups do not utilize profits interests. Likewise, because the partnerships that do utilize profits interests are disproportionately large by dollar value, the top 1% of firms using profits interests account for 98% of the total dollar amount of profits interests.

Third, who is it that enjoys the use of profits interests in terms of the income distribution? [Figure 15](#) in [Appendix B](#) helps flesh out that story. They are highly concentrated, even compared to the distribution of partnership income. Roughly 80% of profits in profits interests flow to the top 1% of earners, compared to only about 70% of partnership profits overall. Carried interests are even more concentrated at the top, with about 90% flowing to the top 1% and roughly half flowing to the top 0.01%.

For purposes of the counterfactual analysis, I estimate that the additional restriction imposed by recharacterizing profits and carried interests as compensation is significant. I find that relative to the no-profits counterfactual, taxes on the same reallocated income are 7-12% higher than under actual allocations across the period of observation, translating to a roughly \$12 billion differential in 2020, and generally \$8 to \$12 billion throughout the ten-year period. This is roughly double the difference under the economic counterfactual. I estimate that the cents of taxes saved per dollar shifted is also higher, steadily at about 10 cents per dollar shifted over the 10-year period.

5.4 Capital Share Counterfactual

The capital share counterfactual imposes substantially more stringent restrictions than the prior two counterfactuals, by disallowing any type of flexible allocation beyond what is analogously possible with S-corporations. Even so, I still find that relatively few partnerships are notably restricted in their behavior. I find that the vast majority of partnerships make allocations consistent with not only their overall economic split between the partners, but also the capital shares of the partners.

I estimate that only about 14% of all partnership allocations are shifted from the economic counterfactual (roughly \$474 billion of total shifting in 2020). But, very similar to the case of the economic counterfactual, this behavior is highly concentrated. I estimate that only 20% of partnership groups shift more than 1% of their allocations, but when they do, they shift on average 35% of their allocations. When measuring by dollar value, the top 1% of shifters

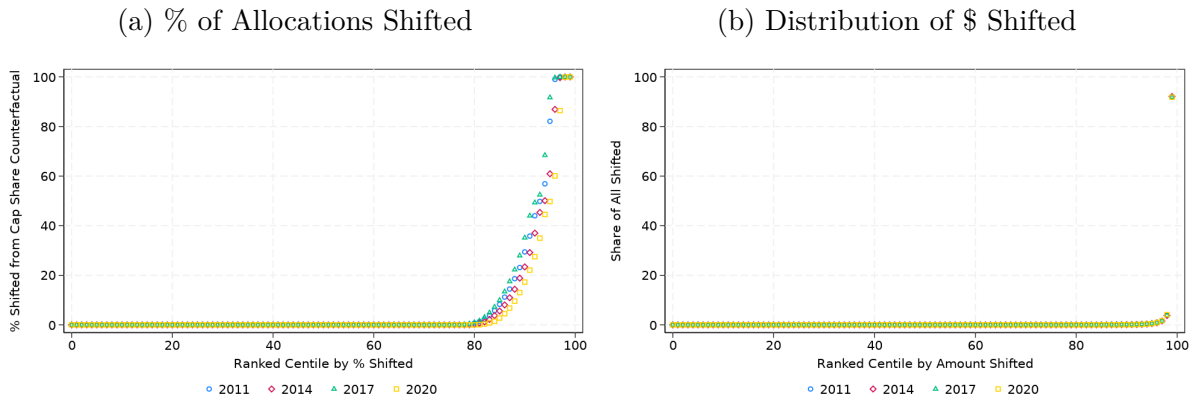
account for about 92% of the total dollars shifted.

The heavy-shifting groups under this counterfactual are similar to those under the economic counterfactual, and this makes sense for two reasons. First, by construction, every partnership that shifts for purposes of the economic counterfactual is also shifting for purposes of the capital share counterfactual. Second, and similarly, the higher shifters under the economic counterfactual tend to also be high shifters under the capital share counterfactual. Groups that shifted more than 1% under the former averaged shifting 38% under the latter. In this regard, it follows that these high-shifter groups under the capital share counterfactual also tend to be larger in both cash flow and number of partners, more likely to be in the Mega Network, more likely to be in the finance industry, more likely to include foreign partners, partners in tax havens, and partners that can't be matched to annual returns. That said, along each of these dimensions, the capital share shifters are much less distinct than the economic counterfactual shifters. In this sense, it appears that the high shifters under the economic counterfactual reflect a particularly rarefied and unique group of partnerships.

Similar to the above discussion, I tend to find other aspects of the capital share shifting group similar to the group of shifters under the economic counterfactual. Similar traits tend to be associated with higher and lower rates of shifting. Similar patterns appear for intensity of shifting across types of income, industry, income percentiles, and firm size. I also see similar patterns for net shifts in allocations across types of partners and income. The one substantial distinction comes from the fact that the capital share counterfactual precludes the use of profits and carried interests, while the economic counterfactual does not. As a result there is much more shifting of ordinary income and capital gains (which are overrepresented in profits and carried interests), and there tends to be very large net shifts of capital gains to individuals (a direct consequence of carried interests).

Unlike the prior counterfactuals, I do not emphasize a tax analysis under this counterfactual. The reason is that the economic interpretation is less intuitive, and is more likely to be misleading. Unlike the prior counterfactuals, the capital share counterfactual causes the economics of the business arrangement to shift. As a result, a change in the tax outcome under this counterfactual may be driven by the fact that completely different partners are receiving the allocations than before, rather than by the mere act of using flexible allocations or profits interests. In light of these points, the best way to interpret a tax analysis under the this counterfactual is to consider it simply a reflection of the extent to which the unique business arrangements facilitated by partnerships are inherently tax-advantageous compared to what might be possible under S-corporations. With these important caveats in mind, I do estimate substantially higher tax liabilities under the capital share counterfactual, even compared to prior results. I estimate

Figure 10: Shifting Under the Capital Share Counterfactual



Notes: The left panel identifies partnership groups (as defined in Section 4), and ranks them (aggregated into centiles) by what percentage of their total allocations are shifted. The right panel shows the distribution of the total amounts shifted (e.g., the point at 99% means that the top centile of groups accounts for 99% of all shifted allocations)

tax liabilities roughly 20% higher than those for actual allocations, translating to about \$23 billion in 2020, and about \$20-30 billion annually during the ten-year period. This translates to about 5 cents per dollar shifted, a lower rate than prior counterfactuals given the much larger amount of dollars shifted.

5.5 Notes on Interpretation of Results

When considering the results above, it is important to keep in mind that both the amount of shifting and the amount of associated tax benefits are likely lower-bound estimates, for several reasons. First, and most importantly, the Schedule K-1 returns I use as the primary data source for present several limitations. They are information returns that do not report tax liabilities, and as such they generally receive less scrutiny and enforcement compared to annual returns like Form 1040 or Form 1120. There are also missing Schedule K-1s, and data errors are not infrequent, as reported by Cooper et al. (2016), Love (2021), and Black et al. (2023). A large number of Schedule K-1s cannot be matched to annual returns of partners, and evidence about these missing returns suggests that these tend to be associated with qualities suggesting more aggressive use of flexible allocations.

Second, not all data is available across all years on the Schedule K-1. Fields excluded as missing or inconsistent for purposes of this project include unrecaptured gains, 1231 gains, exempt income, nondeductible expenses, credits, qualified business income, and all attached schedules and codes. As a result, the quantities of shifted income would be strictly higher.⁴¹

Third, there is extremely limited guidance for reporting capital or profits shares. As a

⁴¹The income I do observe covers about 90% of partnership allocated income, so this loss of information is not especially threatening.

result, a great deal is left to the discretion of the partnership. When there are complex payout waterfalls, it is unclear how shares will be reported. As well, as a conservative manner, where information is missing or difficult to infer, I assume there is no deviation from the capital shares, but this will cause systematic underestimates.

Finally, the magnitudes of the estimates of the tax rate changes are likely attenuated. In many cases, the information available from tax filings is not sufficient to precisely estimate a counterfactual scenario. For example, it is not possible in many cases to identify the sources and amounts of underlying qualified business income for purposes of Section 199A from the information reported, meaning that under a counterfactual scenario the amounts may be either over- or under-estimated. If anything, there is likely greater attenuation among more complex returns, meaning that the high-shifting and high-tax reduction partnerships likely face more attenuation than others.

6 Discussion

My findings contribute to important discussions about tax policy at a time when partnerships are attracting increased attention due to their size, structural complexity, and the inherent challenges of enforcement. As mentioned earlier, partnerships have skyrocketed in terms of their share of businesses and share of profits since the enactment of the check-the-box rules in 1997 (see [Love, 2021](#); [Smith et al., 2022](#); [Black et al., 2023](#)), while the complexity of entity arrangements have also increased as shown here, and with large partnerships growing in number by roughly 600% since 1992.⁴² But yet at the same time audit rates for these large partnerships have plummeted from 1.4% to 0.3%, and the audit rate for partnerships overall has fallen to 0.03%. Moreover, the returns on partnership audits are low, because there has been ineffective targeting and because parsing complex structures is challenging (audits often are not completed in the three year period available) ([GAO, 2023](#)). In light of these trends, plus the fact that partnerships are important components in tax planning, partnerships have been cited as a potentially important part of closing the tax gap ([Sarin and Mazur, 2023](#); [IRS, 2023](#)).

Considering these recent developments, my results can serve to inform policy discussions about partnerships along several dimensions. First, on a technical dimension, the investigation highlights the limits of knowledge from Schedule K-1 and Form 1065 reporting, and how certain additional information could be extremely valuable in mitigating noncompliance with allocation rules. More precisely, one of the main challenges in inferring compliance or noncompliance from current reporting is that that conceptually distinct allocations are commingled: it is impossible,

⁴²Large partnerships are defined by the GAO as those with 100 or more partners and at least \$100 million in assets.

for example, to determine whether a partner is receiving (a) a single large allocation, (b) a small allocation plus a necessary corrective allocation, (c) a small allocation plus a special allocation that is potentially abusive, or some other set of allocations. Separating distinct allocations could significantly ease the challenges currently posed for enforcement in this setting.

On the other hand, it is important to recognize that reporting is burdensome, and extra reporting could strain small businesses. My results inform this aspect of the discussion as well, in that the firms that appear to be shifting the most are not small businesses, but rather large and complex partnerships. But it's also not just size that matters: my findings highlight the importance of partnerships in multi-tier structures, with profits interests, high amounts of portfolio income, and certain entity partners as well. To the extent enhanced reporting may be desired, the findings here help inform the discussion of how to minimize burdens while maintaining efficacy.

Third, the findings here can inform future work aiming to improve the efficiency and usefulness of partnership audits. As reported by the GAO, many audits have been poorly targeted because they are identified using predictors that do not actually predict noncompliance (GAO, 2023). Recent work has expanded upon IRS efforts to identify better predictors (Black et al., 2023), but even so the number of predictor variables evaluated is small and relatively coarse. The descriptive findings of this paper across the many variables evaluated can thus be informative as a basis for future research developing prediction models and audit identification techniques.

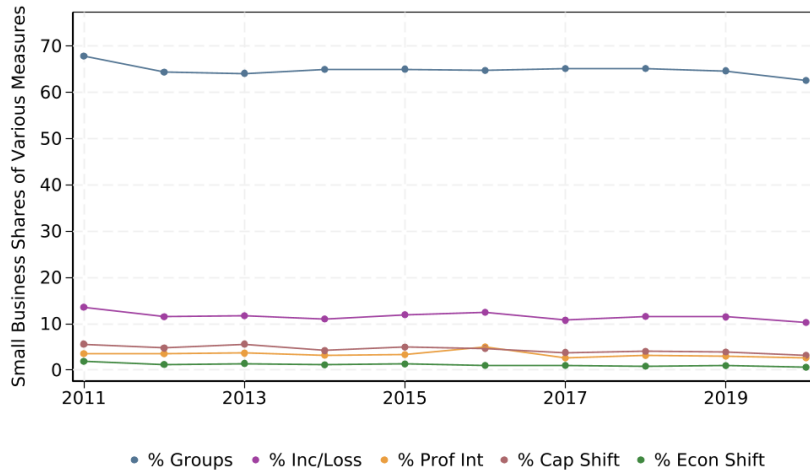
Finally, even if all of the allocations are in compliance with regulations, the findings in this paper raise the deeper and broader question of whether the current partnership regime is delivering intended economic results. One of the motivations for the adoption of the modern partnership tax rules in 1954 was to reduce burdens on businesses, particularly on small businesses. It was observed that among the “great many business enterprises” using the partnership form, it was “commonly employed by small businesses and in farming situations.”⁴³ At the time of the enactment, 61% of partnership profits were attributable to retail, wholesale, and service firms.⁴⁴ But the landscape has since changed dramatically, with 73% now attributable to investment, other finance, and real estate. Moreover, my regression results reveal that it is the larger, more complicated businesses with multi-tier structures, tax haven partners, and portfolio income that are benefiting the most from the use of flexible allocations. This observation raises the question: to what extent are small businesses using these flexible allocations?

In an effort to more precisely measure this concept, I identify partnership groups that one

⁴³S. Rep. No. 1622, 83d Cong., 2d Sess. 89 (1954). See also H.R. Rep. No. 1337, 83d Cong., 2d Sess. 65 (1954).

⁴⁴Data are from 1953, from Business Population, by Legal Status, Series Ch1-292, in *Historical Statistics of the United States, Earliest Times to the Present: Millennial Edition*, edited by Susan B. Carter, Scott Sigmund Gartner, Michael R. Haines, Alan L. Olmstead, Richard Sutch, and Gavin Wright. New York: Cambridge University Press, 2006. <http://dx.doi.org/10.1017/ISBN-9780511132971.Ch1-292>.

Figure 11: Small Businesses as a Portion of All Groups, Income, and Shifting



could reasonably consider “small business” ventures. I categorize groups as small businesses if they satisfy all of the following criteria: (1) less than \$5 million of net income or loss, (2) 50 or fewer partners, all of which are individuals that can be matched to Form 1040s, (3) at most three partnerships in the group, with no circular structures, (4) at least 75% of the partnership income or loss is operating or rental income or loss, (5) no owners are in the top 0.01% of income earners, and (6) the partnership does not use a profits interest to allocate portfolio investment income to partners. Figure 11 shows how these small businesses compare to other partnerships in terms of their counts, income or loss, profits interests, and shifting. First of all, small businesses make up a substantial portion of partnership groups, comprising roughly two-thirds of all partnership groups. But they make up a much smaller portion of total partnership income and loss, roughly 10-15%. They represent an even smaller portion of the total shifting relative to the capital share counterfactual, only constituting about 3-5%. Their share of total profits interests is even small, roughly 2-3%. Finally, their share of shifting under the economic counterfactual is tiny, at only 1%.

7 Conclusion

Partnerships offer more flexibility than any other business entity for owners and investors to divide income and losses. Although this flexibility ideally improves efficiency by giving partners the ability to customize their economic agreement, it also creates incentives to engage in arbitrage, taking advantage of partners’ different tax situations to reduce taxes. This raises the important question: who benefits from this flexibility in practice, and to what extent is it associated with lower taxes?

To investigate this question, I use masked federal tax records from 2011-2020. I take actual

partnership allocations of income and loss and reallocate them by imposing more restrictive sets of rules with less flexibility. I also account for the economics of profits interests and carried interests by compensating service partners with direct payments under the counterfactual (with a corresponding deduction for the partnership). I estimate tax liabilities on the actual versus the counterfactual allocations.

My findings tell a story with several layers. First—perhaps surprisingly—I find that most partnerships do not utilize the flexible allocations unique to partnerships. Only 3% of partnership groups make allocations that “slice-and-dice” individual items of income relative to the overall economic split, only 7% appear to use profits or carried interests, and only 20% use *any* type of flexible allocations unique to partnerships at all.

Second, I find that the partnerships that do engage in such behaviors do so quite intensively, and appear to enjoy significant tax reductions. I estimate that for the partners involved in slicing-and-dicing of individual items there is roughly a 7-10 cent tax reduction per dollar shifted, resulting in a roughly 4% increase in tax liabilities, translating to roughly \$5-7 billion annually. For those that also utilize profits interests and carried interests, the increase is roughly doubled, to about 7-12% or \$8-12 billion.

Who are the beneficiaries of this flexibility? I find that the groups associated with greater use of flexible allocations and greater tax reduction include larger partnerships by cash flow and partner count, more complex entity arrangements with multiple tiers and circular patterns, and partners with disparate tax situations that present opportunities for arbitrage, including partners in tax havens, trusts, and partners with large differences in gain and loss situations. I also find that net shifting appears to flow between types of partners and across types of income in ways that take advantage of partners’ unique tax situations. Taken together, evidence is thus suggestive of tax planning. On the opposite end of the spectrum, I find that small businesses, despite accounting for roughly two-thirds of partnerships, account for a scant portion of this behavior and these benefits.

References

- Agarwal, A., S. Chen, and L. F. Mills (2021). Entity structure and taxes: An analysis of embedded pass-through entities. *The Accounting Review* 96(6), 1–27.
- Black, E., R. Hess, R. Lester, J. Goldin, A. Portz, and D. E. Ho (2023). The spiderweb of partnership tax structures. Technical report, Working Paper.
- Cauble, E. and G. D. Polsky (2014). The problem of abusive related-partner allocations. *Fla. Tax Rev.* 16, 479.
- Collins, J. M. and R. P. Bey (1986). The master limited partnership: an alternative to the corporation. *Financial Management*, 5–14.
- Cooper, M., J. McClelland, J. Pearce, R. Prisinzano, J. Sullivan, D. Yagan, O. Zidar, and E. Zwick (2016). Business in the united states: Who owns it, and how much tax do they pay? *Tax Policy and the Economy* 30(1), 91–128.
- Cuff, T. F. (2009). Several thoughts on drafting target allocation provisions. *Major Tax Plan.* 11- 61, 1.
- Fleischer, V. (2008). Two and twenty: Taxing partnership profits in private equity funds. *NYUL Rev.* 83, 1.
- GAO (2023). Tax enforcement: Irs audit processes can be strengthened to address a growing number of large, complex partnerships. Technical report.
- Gergen, M. P. (1990). Reforming subchapter k: Special allocations. *Tax L. Rev.* 46, 1.
- Gergen, M. P. (1992). Reforming subchapter k: Compensating service partners. *Tax L. Rev.* 48, 69.
- Gergen, M. P. (2003). The end of the revolution in partnership tax. *SMUL Rev.* 56, 343.
- Guyton, J., P. Langetieg, D. Reck, M. Risch, and G. Zucman (2021). Tax evasion at the top of the income distribution: Theory and evidence. Technical report, National Bureau of Economic Research.
- Hasen, D. (2012). Partnership special allocations revisited. *Fla. Tax Rev.* 13, 349.
- IRS (2023). *IRS Inflation Reduction Act Strategic Operating Plan*.
- Knoll, M. S. (2008). The taxation of private equity carried interests: estimating the revenue effects of taxing profit interests as ordinary income. *Wm. & Mary L. Rev.* 50, 115.

- Lokken, L. (2003). As the world of partnership taxation turns. *SMUL Rev.* 56, 365.
- Love, M. (2021). Where in the world does partnership income go? evidence of a growing use of tax havens. *Evidence of a Growing Use of Tax Havens (December 14, 2021)*.
- Monroe, A. (2010). Too big to fail: The problem of partnership allocations. *Va. Tax Rev.* 30, 465.
- Monroe, A. (2012). Integrity in taxation: Rethinking partnership tax. *Ala. L. Rev.* 64, 289.
- Ribstein, L. E. (2009). Partnership governance of large firms. *The University of Chicago Law Review*, 289–309.
- Sarin, N. and M. J. Mazur (2023). The inflation reduction act’s impact on tax compliance—and fiscal sustainability. *Available at SSRN*.
- Smith, M., D. Yagan, O. Zidar, and E. Zwick (2022). The rise of pass-throughs and the decline of the labor share. *American Economic Review: Insights* 4(3), 323–40.
- Yin, G. K. (1997). The taxation of private business enterprises: Some policy questions stimulated by the check-the-box regulations. *SMUL Rev.* 51, 125.

A Binary Variables in Regression

Table 3: Binary Variables Used in Regressions, Part 1

Variable	Description
Size Variables	
· Large Partnership	Group has large partnership (100 partners + \$100M assets, per GAO)
· >\$500m Inc/Loss	Group has more than \$500M income or loss
· \$5-500m Inc/Loss	Group has between \$5M and \$500M income or loss
· <\$5m Inc/Loss	Group has less than \$5M income or loss
· >\$500m Assets	Group has more than \$500M assets
· <\$1m Assets	Group has less than \$1M assets
· Small Business	Group is a “small business” as defined in text
Network Variables	
· 2 Partners	Group has exactly 2 non-partnership partners
· <10 Partners	Group has fewer than 10 non-partnership partners
· 100+ Partners	Group has at least 100 non-partnership partners
· 4+ Pshp Group	Group has at least 4 partnerships
· Circular < 4 Pshp	Group has partnerships part of a small circular structure (< 4 pshps)
· Circular 50+ Pshp	Group has partnerships part of a large circular structure (≥ 50 pshps)
· Solo Partnership Network	Network has only 1 partnership
· Mega Network	Group is part of the Mega Network
· Only 1 Tier	Group has only 1 tier and is not above any other partnerships
· 4+ Tiers	Group has partnerships that are above ≥ 4 tiers of lower partnerships
· Above Investment Fund	Group either is investment fund or is above a lower-tier fund
Entity Variables	
· All Individual Partners	All non-partnership partners are individuals
· All Entity Partners	All non-partnership partners are entities
· Trust Partner	Group has a trust as a partner
· 25%+ to Trusts	At least 25% of income/loss goes to trusts
· S-corp Partner	Group has an S-corp partner
· 25%+ to S-corps	At least 25% of income/loss goes to S-corp
· 10%+ to Unmatched K-1s	At least 10% of income/loss goes to K-1s that aren’t matched to returns
· 25%+ to Unmatched K-1s	At least 25% of income/loss goes to K-1s that aren’t matched to returns
· 50%+ to C-corps	At least 50% of income/loss goes to C-corps
· 50%+ to Partnerships	At least 50% of income/loss goes to partnerships

Notes: All variables take a value of 1 or 0 at the Group level (observations are at the group level) based upon whether or not the characteristic applies.

Table 5: Binary Variables Used in Regressions, Part 2

Variable	Description
K-1 Income Variables	
· Mostly Operating Income	At least 75% of income/loss is operating
· Mostly Rental Income	At least 75% of income/loss is rental
· Mostly Interest Income	At least 75% of income/loss is interest
· Mostly Cap Gain	At least 75% of income/loss is capital gains
· Mostly Portfolio Income	At least 75% of income/loss is portfolio income
Industry Variables	
· 50%+ Finance	At least 50% of income/loss is in Finance industry
· 50%+ Real Estate	At least 50% of income/loss is in Real Estate industry
· 50%+ Oil & Gas	At least 50% of income/loss is in Oil & Gas industry
· 50%+ Manufacturing	At least 50% of income/loss is in Manufacturing industry
· 50%+ Trade	At least 50% of income/loss is in Retail or Wholesale industry
· 50%+ Services	At least 50% of income/loss is in Prof., Health, or Education services
· Has Investment Fund	Group includes an investment fund
Geography Variables	
· 25%+ Tax Haven	At least 25% of income/loss goes to a haven
· 25%+ Cayman	At least 25% of income/loss goes to Cayman Islands
· 25%+ to Foreign	At least 25% of income/loss goes to foreign partners
Distribution Variables	
· \$\$\$ to Top 0.01 Percentile	At least 25% of income/loss goes to individuals in top 0.01%
· \$\$\$ to Top 1 Percentile	At least 25% of income/loss goes to individuals in top 1%
Capital Share Variables	
· High Cap Share Variation	Top 25% in variation of capital share between partners
· Low Cap Share Variation	Bottom 50% in variation of capital share between partners
Profits Interest Variables	
· > 10% Profits Interest	More than 10% of all profits go to a profits interest partner.
· Large Profits Interest	A partner has a profits interest of at least 10%.
Allocation Variables	
· High Allocation Variation	Top 25% in variation of K-1 allocations between partners
· Low Allocation Variation	Bottom 25% in variation of K-1 allocations between partners
Tax Situation Variables	
· High Partner Tax Variation	Top 25% in variation of outside income/loss to partners
· Low Partner Tax Variation	Bottom 25% in variation of outside income/loss to partners
· \$\$\$ to Gain & Loss Partners	At least 25% of Inc/Loss going to loss partners and gain partners
· High Partner Outside Income	Group is top 25% for outside net income of partners
· High Partner Outside Loss	Group is top 10% for outside net loss of partners
· Top 25% Ordinary Rates	Group is top 25% for marginal ordinary tax rate of partners
· High Var Ordinary Rates	Group is top 25% for variation in marginal ordinary tax rate of partners
· Top 25% Cap Gain Rates	Group is top 25% for marginal cap gains tax rate of partners
· High Var Cap Gain Rates	Group is top 25% for variation in marginal cap gains tax rate of partners
Shifting Variables	
· Top 5% Shifted	Group is top 5% in terms of share of income/loss shifted
Lag Variables	
· Lag >1pp Response	Group last year had >1pp response

Notes: All variables take a value of 1 or 0 at the Group level (observations are at the group level) based upon whether or not the characteristic applies.

Table 7: Means of Binary Variables

Variable	Unweighted		Inc/Loss Weighted	
	Mean	SE	Mean	SE
Large Partnership	0.0008	0.0000	0.1531	0.0001
>\$500m Inc/Loss	0.0003	0.0000	0.0889	0.0001
>\$500m Assets	0.0018	0.0000	0.2824	0.0001
<\$1m Assets	0.7652	0.0001	0.3266	0.0001
\$5-500m Inc/Loss	0.0011	0.0000	0.0643	0.0000
<\$5m Inc/Loss	0.9949	0.0000	0.7907	0.0001
Small Business	0.6477	0.0001	0.1592	0.0001
2 Partners	0.9405	0.0000	0.5493	0.0001
100+ Partners	0.0027	0.0000	0.1737	0.0001
4+ Pshp Group	0.0027	0.0000	0.0852	0.0001
50+ Pshp Group	0.0000	0.0000	0.0083	0.0000
Circular <4 Pshp	0.0003	0.0000	0.0016	0.0000
Circular 50+ Pshp	0.0002	0.0000	0.0171	0.0000
Solo Partnership Network	0.8475	0.0001	0.4307	0.0001
Mega Network	0.0905	0.0001	0.4726	0.0001
Only 1 Tier	0.9828	0.0000	0.8214	0.0001
4+ Tiers	0.0002	0.0000	0.0151	0.0000
Above Investment Fund	0.0554	0.0000	0.3661	0.0001
All Individual Partners	0.6996	0.0001	0.2136	0.0001
All Entity Partners	0.0596	0.0000	0.1381	0.0001
Trust Partner	0.1202	0.0001	0.4840	0.0001
25%+ to Trusts	0.0781	0.0001	0.1808	0.0001
S-corp Partner	0.0682	0.0000	0.2679	0.0001
25%+ to S-corps	0.0338	0.0000	0.0653	0.0000
25%+ to Unmatched K-1s	0.1224	0.0001	0.2610	0.0001
50%+ C-corps	0.0292	0.0000	0.1686	0.0001
50%+ Partnerships	0.0776	0.0001	0.1184	0.0001
50%+ Finance	0.1708	0.0001	0.4746	0.0001
50%+ Real Estate	0.3544	0.0001	0.1468	0.0001
50%+ Oil & Gas	0.0025	0.0000	0.0127	0.0000
50%+ Manufacturing	0.0166	0.0000	0.0314	0.0000
50%+ Trade	0.0616	0.0000	0.0310	0.0000
50%+ Services	0.0817	0.0001	0.1157	0.0001
Has Investment Fund	0.1644	0.0001	0.4742	0.0001
25%+ Tax Haven	0.0049	0.0000	0.0616	0.0000
25%+ Cayman	0.0019	0.0000	0.0350	0.0000
25%+ to Foreign	0.0472	0.0000	0.1157	0.0001
Mostly Operating Income	0.5206	0.0001	0.4106	0.0001
Mostly Rental Income	0.4313	0.0001	0.1262	0.0001
Mostly Interest	0.0874	0.0001	0.0229	0.0000
Mostly Cap Gain	0.0966	0.0001	0.2275	0.0001
Mostly Portfolio Income	0.1419	0.0001	0.3678	0.0001
\$\$\$ to Top 0.01 Percentile	0.0197	0.0000	0.2697	0.0001
\$\$\$ to Top 1 Percentile	0.2914	0.0001	0.7011	0.0001
High Cap Share Variation	0.2489	0.0001	0.2740	0.0001
Low Cap Share Variation	0.4977	0.0001	0.2063	0.0001
>10% in Profits Interest	0.0350	0.0000	0.1135	0.0001
Large Profits Interest	0.0268	0.0000	0.1066	0.0001
High Allocation Variation	0.2489	0.0001	0.8823	0.0001
Low Allocation Variation	0.2489	0.0001	0.0328	0.0000
High Partner Tax Variation	0.2287	0.0001	0.7990	0.0001
Low Partner Tax Variation	0.2287	0.0001	0.0335	0.0000
\$\$\$ to Gain & Loss Partners	0.0129	0.0000	0.0522	0.0000
High Partner Outside Inc	0.2500	0.0001	0.2413	0.0001
High Partner Outside Loss	0.1000	0.0001	0.6499	0.0001
Top 25% Ordinary Rates	0.2500	0.0001	0.1599	0.0001
Top 25% Cap Gain Rates	0.2500	0.0001	0.2405	0.0001
High Var Ord Rates	0.2414	0.0001	0.4304	0.0001
High Var Cap Gain Rates	0.2414	0.0001	0.3042	0.0001
Top 5% Shifted (Economic)	0.0464	0.0000	0.1467	0.0001

B Additional Figures

Figure 12: Shifting Intensity, Economic Counterfactual

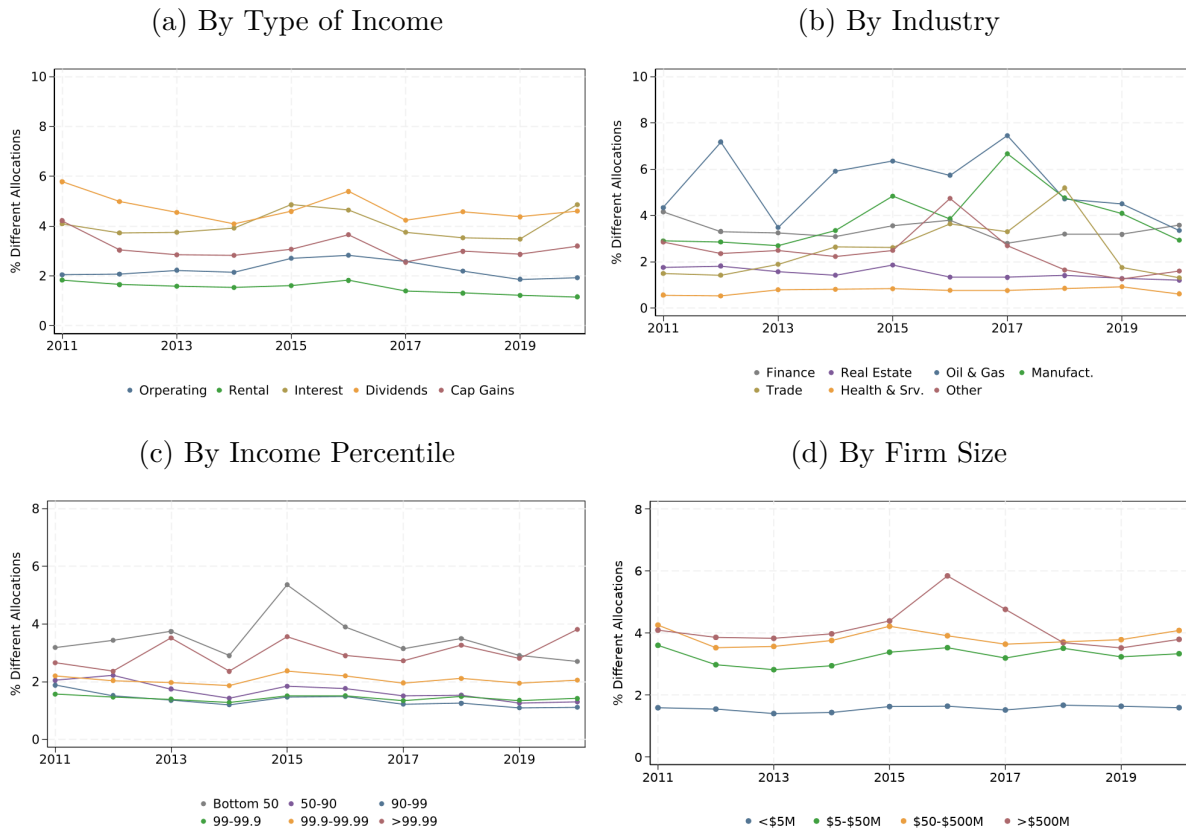
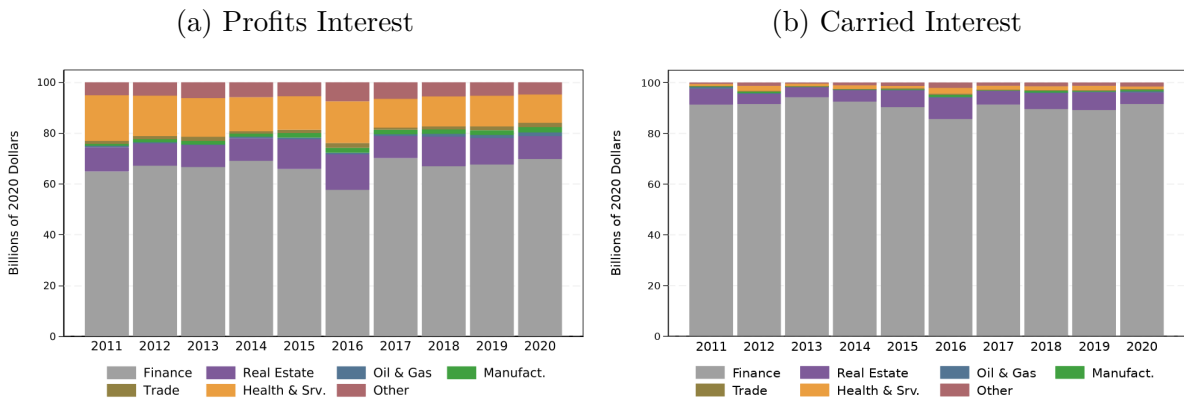


Figure 13: Shifting Intensity, Capital Share Counterfactual



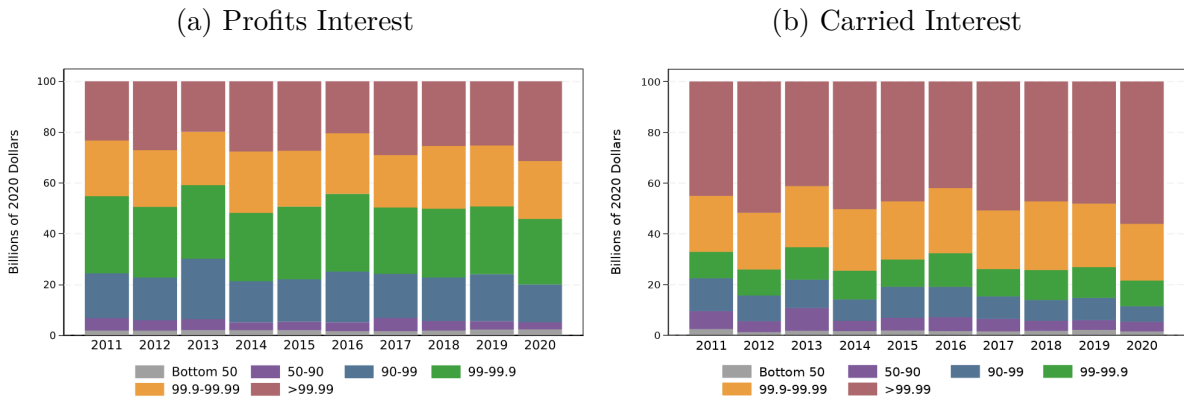
Notes: Each panel shows the shifting intensity by category, over time, measured as the percentage of allocations shifted.

Figure 14: Profits & Carried Interests by Industry




Notes: The left panel shows the industry breakdown of the partnerships issuing profits interests, weighted by amount of profits passing through the profits interest. The right does the same for carried interest.

Figure 15: Profits & Carried Interests by Income Percentile




Notes: The left panel shows the distribution of the receipt of profits from profits interests across the income distribution, weighted by amount of profits received. Profits interests are only included in this measure to the extent they could be matched to form 1040 filings. The right panel does the same for carried interest.

Figure 16: Profits Interests by Type of Income



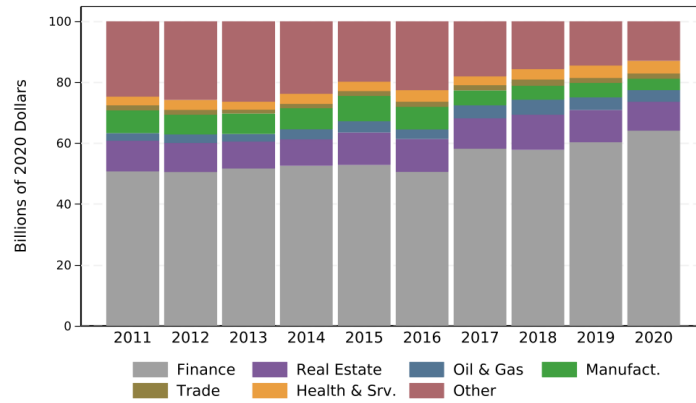
Notes: This figure the breakdown of profits interests by type of income received.

Figure 17: Missing Partners by Entity Type




Notes: This figure presents the distribution of profits and losses flowing to missing partners, broken down by entity type.

Figure 18: Missing Partners by Industry




Notes: This figure presents the distribution of profits and losses flowing to missing partners, broken down by industry.

Figure 19: Missing Partners by Type of Income



Notes: This figure presents the distribution of profits and losses flowing to missing partners, broken down by type of income.

Figure 20: Missing Partners by Group Size



Notes: This figure presents the distribution of profits and losses flowing to missing partners, broken down by group size.

Figure 21: Income Distribution: Solo Networks vs. Mega Network

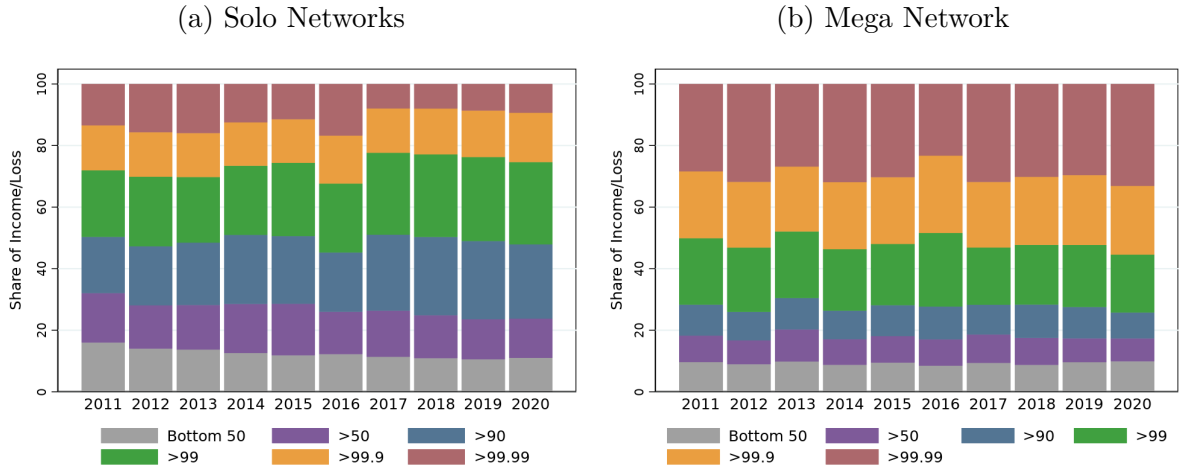


Figure 22: Income Type: Solo Networks vs. Mega Network

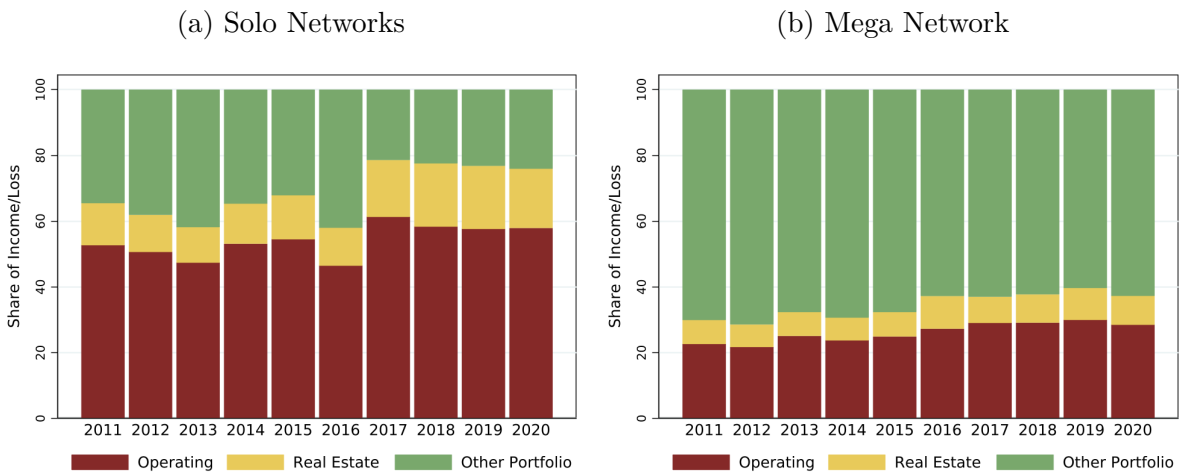



Figure 23: Group Size: Solo Networks vs. Mega Network

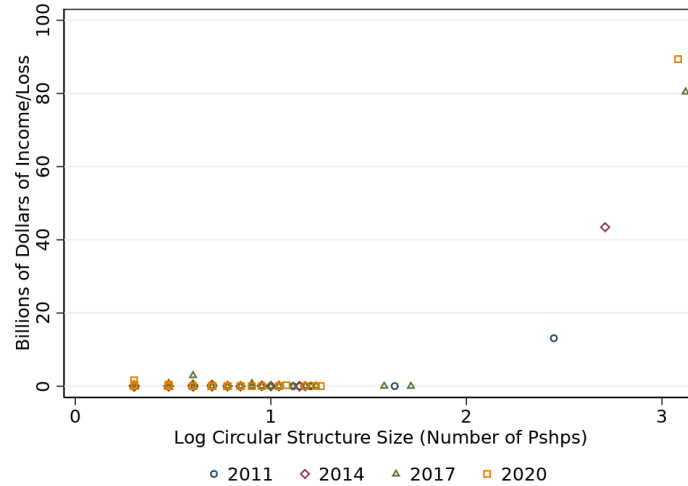
(a) Solo Networks

(b) Mega Network




Notes: Each bar represents the composition of the respective networks in terms of total dollars of income and loss reported on Schedule K-1s for a given year.

Figure 24: Most Circular Income/Loss is in a Single Big Structure



Notes: This figure identifies circular structures, and measures the total income and loss reported in each, as well as the number of partnerships in each structure. The relationship between size of the structures in terms of partnership counts and in terms of income and losses is plotted, across time.

Figure 25: Use of Flexible Allocations Over Time



Notes: This figure identifies the amount shifted when comparing actual allocations versus the two counterfactuals, calculated according to Equation 1, over time.

Table 8: Persistence of Shifting Over Years (Econ)

Group	Original Share	New Group	+1 Year	+2 Years	+3 Years
<1%	97.8%	Closed	12%	21%	28%
		<1%	87%	78%	71%
		1-50%	1%	1%	1%
		>50%	0%	0%	0%
1-50%	2.0%	Closed	18%	26%	32%
		<1%	48%	47%	44%
		1-50%	32%	26%	22%
		>50%	2%	2%	2%
>50%	0.2%	Closed	21%	30%	37%
		<1%	50%	47%	44%
		1-50%	20%	17%	14%
		>50%	9%	6%	5%

Notes: The table describes the initial distribution in a given year across different thresholds of shifting, and then tracks these firms using the panel data to see how shifting persists over time. The percentages in each subsequent year describe how firms in each group move to new groups or stay in their original group over time.

Table 9: Persistence of Shifting Over Years (Cap Share)

Group	Original Share	New Group	+1 Year	+2 Years	+3 Years
<1%	79.5%	Closed	12%	20%	27%
		<1%	81%	72%	66%
		1-50%	4%	4%	4%
		>50%	3%	3%	3%
1-50%	10.7%	Closed	12%	20%	28%
		<1%	30%	30%	29%
		1-50%	50%	42%	36%
		>50%	8%	8%	7%
>50%	9.8%	Closed	16%	27%	35%
		<1%	33%	30%	29%
		1-50%	11%	10%	10%
		>50%	41%	32%	26%

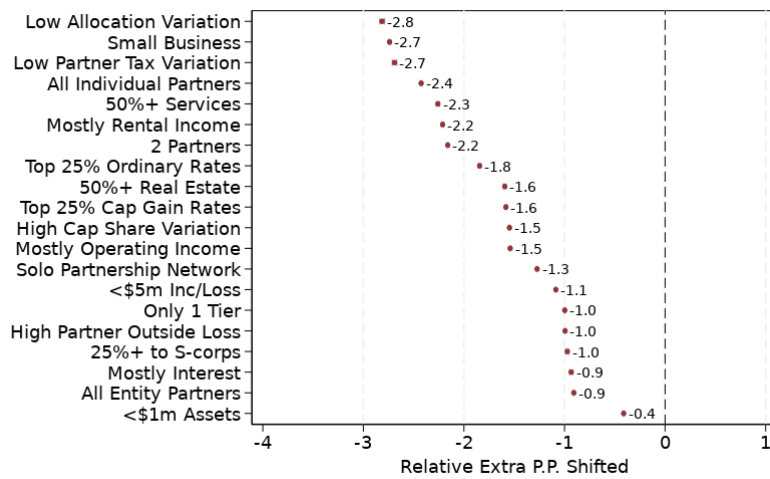
Notes: The table describes the initial distribution in a given year across different thresholds of shifting, and then tracks these firms using the panel data to see how shifting persists over time. The percentages in each subsequent year describe how firms in each group move to new groups or stay in their original group over time.

Figure 26: Net Shifting To and From Types of Partners, by Income Type (Econ)

	Ordinary		Rental		Interest		Dividends		ST Cap		LT Cap	
	+	-	+	-	+	-	+	-	+	-	+	-
Individual	62	-67	13	-13	15	-19	14	-18	30	-27	59	-47
Grantor Trust	8	-14	2	-2	5	-4	5	-5	8	-9	15	-13
C-Corp	28	-33	4	-4	4	-5	2	-3	2	-2	7	-8
S-Corp	8	-7	1	-1	1	-1	1	-1	2	-1	3	-3
1041 Trust	5	-7	2	-2	3	-3	4	-4	6	-6	9	-9
Tax-Exempt	6	-7	1	-1	7	-9	12	-10	12	-13	21	-24
Unmatched Corporation	47	-34	5	-6	7	-6	4	-5	5	-5	11	-11
Unmatched Passthrough	20	-13	3	-3	5	-3	4	-3	4	-4	8	-10
Unmatched Entity	10	-9	2	-2	2	-3	3	-3	5	-4	5	-5
Other US	12	-12	5	-5	4	-3	5	-4	5	-5	8	-9
1120-F	4	-5	1	-1	5	-3	4	-3	3	-4	7	-9
Other Foreign	13	-17	4	-1	9	-7	10	-9	11	-12	17	-21

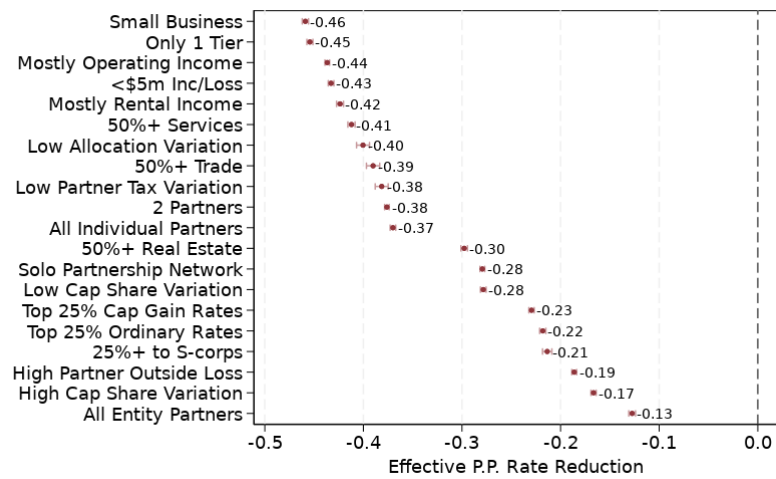
Notes: The table reports the net amounts shifted to (+) or from (-) each type of partner, by type of income, reported in billions of dollars.

Figure 27: Low Shifting from Economic Counterfactual, by Group Trait




Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the share shifted at the group level, following Equation 1. Observations are at the group level. Point estimates represent p.p. increases in the shifting rate associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with least shifting.

Figure 28: Rate Reduction Under Economic Counterfactual, by Group Trait (Low)



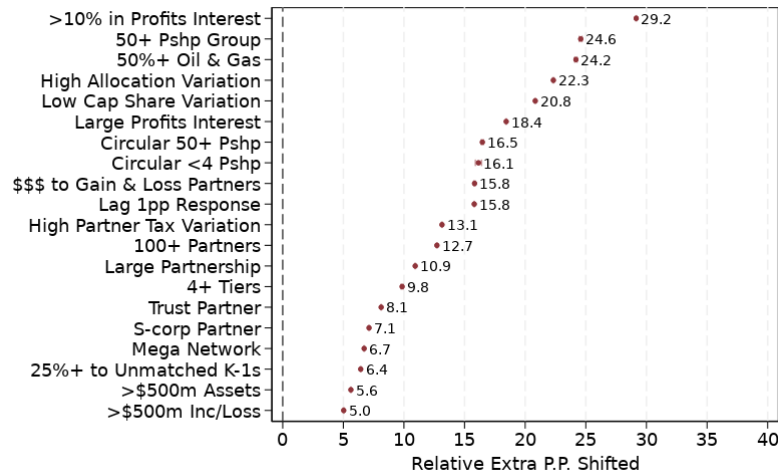
Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the percentage point rate reduction at the group level. Observations are at the group level. Point estimates represent p.p. reduction in the effective tax rate associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with the lowest net reductions.

Figure 29: Cents of Tax Reduction per Dollar Shifted, Economic Counterfactual (Low)



Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the cents of tax reduction per dollar shifted at the group level. Observations are at the group level. Point estimates represent p.p. reduction associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with lowest net reductions.

Figure 30: High Shifting from Capital Share Counterfactual, by Group Trait



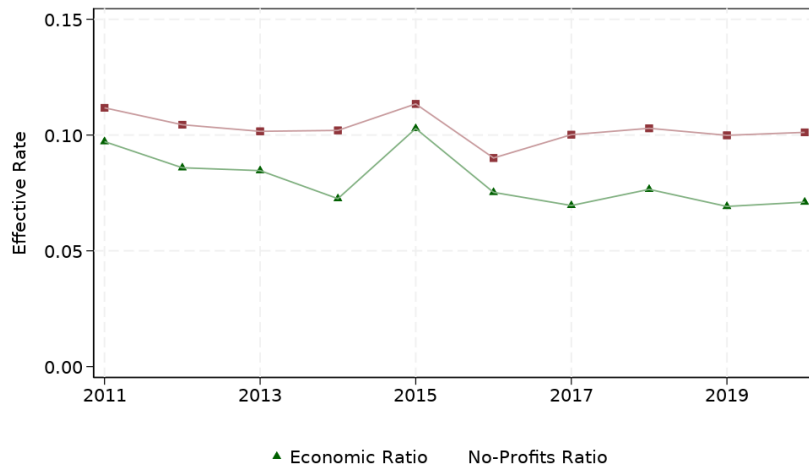
Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the share shifted at the group level, following Equation 1. Observations are at the group level. Point estimates represent p.p. increases in the shifting rate associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with most shifting.

Figure 31: Low Shifting from Capital Share Counterfactual, by Group Trait



Notes: This figure shows the results of separate regressions running each dummy as a sole predictor variable. The dependent variable is the share shifted at the group level, following Equation 1. Observations are at the group level. Point estimates represent p.p. increases in the shifting rate associated with the trait, relative to groups without the trait. Error bars represent 95% confidence intervals. Observations are weighted by the absolute value of the total income and loss reported by the group. Year fixed effects are included. This plot represents the 20 traits associated with least shifting.

Figure 32: Cents of Tax Reduction per Dollar Shifted



Notes: Numerator is the net tax reduction under each counterfactual, denominator is the total amount shifted.