

DANIEL N. SHAVIRO

FIXING U.S. INTERNATIONAL TAXATION

OXFORD

FIXING U.S. INTERNATIONAL TAXATION

(forthcoming Oxford University Press, 2014)

Daniel N. Shaviro

June 2013

All Rights Reserved

1. INTRODUCTION AND OVERVIEW

A Fork in the Road?

Yogi Berra once offered the advice, “When you come to a fork in the road, take it.” He could almost have been speaking about the U.S. international tax rules, which govern how we tax cross-border or multinational investment. For “outbound” investment, or that earned abroad by U.S. companies, the U.S. rules, for almost a century, have muddled along in the netherworld between two sharply etched approaches that dominate the literature, each intuitively appealing but utterly inconsistent with the other.

The first approach is called *worldwide* or *residence-based* taxation. Under it, the U.S. would impose tax at the same rate on U.S. companies’ foreign source income (FSI) as on their domestic income. The great apparent virtue of this approach is that it would prevent the companies from reducing their U.S. tax liability by investing (or reporting income) abroad rather than at home.

The second approach is called *source-based* or *territorial* taxation. Under it, the U.S., recognizing that foreign companies pay no U.S. tax when they invest abroad, would extend this same U.S. tax exemption for FSI to its own companies. (This approach is therefore also called *exemption*). The great apparent virtue of this approach is that it would avoid placing U.S. companies under a competitive disadvantage, as compared to their foreign rivals, when they invest abroad.

How could one possibly object to either approach? Surely no one wants the U.S. tax rules to encourage our companies to invest (or report income) abroad rather than at home. That would seemingly be self-defeating. Yet surely no one wants U.S. firms to face competitive handicaps when they invest abroad. Whether one favors a “level

playing field” – long one of tax policy’s most powerful and popular (if trite) metaphors – or is actively rooting for “our” home team firms to “win” their contests on the road, imposing a tax handicap on U.S. firms seems self-defeating as well.

Unfortunately, however, we cannot follow both of these approaches at once. For example, if we tax General Electric (GE) when it invests in the U.S., and do not tax GE’s global rival Siemens when it invests in Germany, we simply cannot tax GE’s German income both at the U.S. domestic rate and at our zero tax rate for Siemens’ German income. Rather, we must choose one or the other, or else compromise with something in the middle that fully satisfies neither approach.

In stating the dilemma this way, I have diverged from the prevailing custom in international tax policy debate, which is to compare *all* countries’ taxes on a given investment, rather than just those imposed by a given country (such as the United States). My reasons for describing the problem this way, rather than in the conventional way, will become clear later in this chapter and throughout the book. But even if one takes account of all countries’ taxes when checking for tax bias, the same problem arises unless everyone charges the same rate. Thus, suppose the U.S. tax rate is 35 percent, while that in Germany is only 20 percent. In these circumstances, the U.S. would be unable to match GE’s overall global tax rate when it invested in Germany both to the 35 percent rate that it would pay domestically, and to the 20 percent rate that Siemens was paying in Germany.

Accordingly, whether one looks just at U.S. taxes or at all countries’ taxes (where their rates differ), one faces the same dilemma. Under either formulation, concern about GE’s tax minimization incentives suggests imposing U.S. tax (or at least a make-up tax)

on its German income, while concern about the GE vs. Siemens differential suggests that the U.S. exempt non-U.S. source income. Yet, however intuitively appealing we might find each approach, we cannot follow them both.

Reflecting the evident awkwardness of this dilemma, the U.S. rules, for almost a century, have been suspended in the middle between the worldwide and territorial poles, taxing U.S. companies' FSI, but at a greatly reduced effective rate compared to that which applies domestically. Opposing political forces, backed by the rival approaches' offsetting intuitive appeal, have long battled to a stalemate that not only satisfies neither side, but is absurdly dysfunctional due to its elaborate Rube Goldberg-style structure. Not merely unaesthetic (though surely that as well), this structure makes it inevitable that the planning and compliance costs the system induces will be "disproportionately high relative to their role in the activities of the corporation," and "extremely high relative to the revenue raised by the U.S. government on this income" (Blumenthal and Slemrod 1996, 48).

The years since this statement was made have only added to its truth, due to the development of all sorts of new tax planning tricks that empower U.S. companies to lower their U.S. (and foreign) tax obligations at the cost of extra paperwork and convoluted internal cash flow management. Making things worse still, all of this "costly tax planning ... necessitate[s] a significant amount of costly enforcement and compliance activities by the IRS" (President's Economic Recovery and Advisory Board 2010, 88). Given all this, "[m]ost experts agree that the current hybrid U.S. system ...embodies the worst features of both a pure worldwide system and a pure territorial system from the perspective of simplicity, enforcement, and compliance" (88).

Imposing extremely high tax planning and compliance costs relative to the revenue raised is a strong indicator of inefficiency. However, the problems with the existing compromise system for taxing U.S. companies' FSI do not end here. Supporters of worldwide residence-based U.S. corporate taxation complain that the current system encourages rampant tax avoidance by U.S. multinationals, undermining broader tax equity and causing the companies to be unduly tax-favored relative to their purely domestic competitors. Supporters of exempting FSI complain, not only about the potential disadvantages of being a U.S. rather than a foreign multinational when one invests abroad, but also about the problem of "trapped" foreign earnings under our current rules, which may stay abroad for U.S. tax reasons even if they otherwise would come straight home to fund U.S. investments or be distributed to U.S. shareholders. Despite the seeming tension between these two accounts, both are largely correct.

Unfortunately, near-universal consensus that the existing U.S. international tax system is horrendously bad has failed to induce change, given the continuing dissensus about what to replace it with. An obvious solution might involve moving decisively to one of the two poles, by more fully adopting some version of either worldwide taxation or territoriality. This, however, has been impeded by the fact that, in politics no less than trench warfare, it is easier to defend one's own position than to advance through the enemy's.

Of late, there have been some signs of a possible shift towards territoriality. Long a Republican goal, there were occasional hints during the first Obama Administration that Democrats might embrace it as well.¹ Factors encouraging the shift, beyond a generally pro-business and anti-tax political environment, include the rising pressures of global tax

competition, and the fact that peer countries, such as the United Kingdom and Japan, have recently made their international tax systems more territorial. As a political matter, it is plausible that the U.S. would already have shifted to a territorial system if not for budgetary concerns. A recent Treasury estimate found that enacting a simplified territorial system without safeguards might reduce U.S. revenues by \$130 billion over ten years (President's Economic Recovery Advisory Board 2010, 90). Moreover, while a package of international tax law changes that included exemption could, on the whole, pay for itself or even raise revenue, depending on the other items in the package, this might kill all the fun, so far as exemption's main political supporters are concerned.

Among the factors that have boosted exemption's political prospects is a shift in the intellectual climate of debate among U.S. international tax policy experts. Worldwide residence-based taxation used to be considered clearly superior by most experts – better for the world, and better for the United States. However, recent theoretical and empirical work has destroyed this consensus, and for a while seemingly raised the prospect of an emerging opposite consensus, in which a purely territorial tax was seen as better both for the world and for us. The main concern that continues to get in the way has less to do with taxing FSI as an end in itself, than with the view that worldwide taxation is indispensable to combating profit-shifting by U.S. companies, which have become ever more adept at treating large shares of their U.S. (and other) profits as having arisen in tax havens for official reporting purposes (see, e.g., Kleinbard 2011a).

The Need for a New Framework

No matter how the short-term politics plays out – continued stalemate that might perhaps include tougher limits on multinationals' tax planning opportunities, or shift

towards exemption – ongoing disputes concerning how the U.S. ought to tax domestic and foreign multinationals are unlikely to recede any time soon. Consider the episodes of media excitement in recent years over aggressive tax planning by such iconic U.S. companies as Apple, GE, and Google. The continuing rise of both global economic integration and tax planning technology, along with ongoing anxieties about U.S. job levels that both sides in the U.S. international tax debate invoke as favoring their preferred approaches, will keep these issues prominent indefinitely. Thus, it would be desirable to have a clear intellectual framework on hand. Best of all, of course, would be to find consensus “right answers” to all of the hard issues in the field. Even short of that, however, there is a need for better guidance regarding to how to think about the main tradeoffs.

Making this more difficult is the fact that the international tax policy literature, despite a more than fifty year history of frequently intensive academic study by exceptionally talented and knowledgeable economists and lawyers, at some point went badly off the rails. For many decades, its main terms of debate have too often reflected crucial misunderstandings of key issues and distinctions, along with a misguided focus on concepts that verge on being completely unhelpful.² We need therefore, to start again from first principles – albeit, principles that are routinely used elsewhere in public economics. To begin in demonstrating the need for a new approach, consider the following five conceptual problems in the existing international tax policy literature:

First Problem: What *Is* the Problem? Is It Double (and Non-) Taxation?

Popular and academic writers on international taxation – albeit lawyers more than economists – are largely at one in defining the basic problem that the rules in the field

must try to address. It supposedly is “double taxation,” which may arise when cross-border investment faces both a residence-based tax in the home country and a source-based tax in the country where the investment is made. Preventing double taxation from occurring – along with double non-taxation, which occurs when cross-border income is taxed nowhere – is thus identified as the supreme task of international tax rules.

To illustrate what “double taxation” is taken to mean, suppose that Acme Products, a U.S. company, earns \$100 in Germany, and that the U.S. rate is 35 percent, while the German rate is 20 percent. Taking it as given that Germany will in fact charge Acme \$20 of tax, the U.S. is thought to have only two legitimate alternatives. The first is to apply exemption, so that Acme’s German income is only taxed in Germany. The second is to charge U.S. tax – it always is assumed, at the same 35 percent statutory rate that applies domestically – but to provide foreign tax credits, or a dollar-for-dollar reduction of the U.S. tax due by reason of the German taxes paid.

Under this second approach, the \$35 of U.S. tax liability that Acme would otherwise have owed on its (pre-German tax) German income is reduced to \$15 by reason of the allowable credits. Thus, while Acme is literally paying tax to two countries, it is not viewed as being “double-taxed,” because the U.S. credits entirely wash out the burden on Acme of having to pay the German tax. Thus, it is just as if Acme had only been taxed by the U.S., but with a side payment from the U.S. to the German Treasury wholly compensating the latter for its forbearance.

Accordingly, the U.S. tax on Acme’s German income ostensibly must be either zero or \$15, as opposed to anything in between. This alone should excite suspicion. One would normally expect policy assessments to exhibit greater continuity. Given, for

example, that raising the tax from zero to \$1, or lowering it from \$15 to \$14, seems unlikely to change dramatically its practical effects (such as on the central tax policy concerns of efficiency and distribution), it is surprising to witness the evident assumption that everything in the middle must be ruled out.

Why is aversion to double taxation – leading to the assumption that the U.S. must either exempt foreign source income or offer foreign tax credits – so prevalent? The reasons appear to be several. First, there evidently is some sort of inchoate moral intuition lying behind it – although, as we will see, the intuition proves upon closer examination to be incoherent and formalistic.

Second, bilateral tax treaties, such as the existing one between the U.S. and Germany, invariably use the concept of avoiding double (and non-) taxation as a coordinating device between overlapping tax systems. As we will see, in the bilateral setting of a treaty, this concept may indeed be convenient and useful. But the U.S. is not operating in the treaty setting when it more generally determines how it should tax outbound investment.

The third ground for treating aversion to double taxation as a fundamental principle should help make its misdirectedness – as well as its relationship to what actually are serious concerns – somewhat clearer. Suppose, in the prior example, that the U.S. taxed Acme's \$100 of German income at the full 35 percent rate, with the German taxes of \$20 merely being deductible. Then Acme would end with only \$52 after paying both German and U.S. taxes. Its overall global tax rate, on this income, of 48 percent would compare quite unfavorably with the global rates (i.e., just the domestic rates) that U.S. and German companies would face when investing at home. And German

companies, of course, would face the same 48 percent global tax rate when they invested in the U.S., if Germany taxed its residents' foreign source income at the full domestic rate, with U.S. taxes merely being deductible. The result would likely be significant discouragement of cross-border investment by U.S. and German firms.

The prospect of such discouragement is indeed relevant to U.S. policymaking, whether we are setting our international tax policy rules unilaterally or engaged in strategic interaction with other countries. However, the use of such examples to support an anti-double tax principle confuses two distinct questions. The first is "How much?" The second is "How many times?"

To illustrate the distinction between these two questions, suppose the U.S. treated German taxes as merely deductible, but lowered its tax rate for FSI to 15 percent. Acme would now be getting double-taxed, since the U.S. would no longer be entirely negating the separate impact of the German tax. But Acme's U.S. tax bill on its FSI would have declined from \$15 under the full rate / foreign tax credit system, to just \$12 (i.e., 15 percent of \$80). So Acme's outbound investment would actually be *less* tax-discouraged than previously, and the "crime" of double-taxation would evidently be victimless.

Even Germany, while it could point to a clear treaty violation, would face no adverse consequence apart from one that actually is unrelated to double taxation as such. The change in U.S. policy might indeed make Germany worse-off, since U.S. firms that were effectively subject to the U.S. worldwide tax would no longer be able to credit their German taxes against those levied by the U.S. Thus, U.S. multinationals might be more eager than previously to minimize their German tax liabilities, whether through tax planning or actual shifts in where they operate. But they would likewise have this

incentive if the U.S., consistently with its tax treaty obligations, simply exempted those companies' FSI from the prospect of facing any residual U.S. tax. Indeed, under exemption they would care *more* about avoiding German taxes than under a worldwide / foreign tax deductibility approach with, say, a 15 percent U.S. tax rate on FSI, since the latter would reduce their after-U.S. tax cost of paying a dollar of German taxes from one full dollar to 85 cents.

This example helps to show that exemption is what I call an implicit deductibility system. That is, it has the same effect on U.S. companies' incentives, in trading off foreign tax liabilities against other foreign expenses (or forgone income), as would the adoption of explicit foreign tax deductibility with a positive U.S. tax rate (whether, say, 0.00001 percent or 35 percent) for foreign source income.

In sum, while there may be good reason to care about high versus low taxes on cross-border activity, or about even versus uneven taxes at a given margin, this does not imply that there is any direct normative reason to care about the *number* of taxes that are being levied on a given taxpayer or transaction. After all, most of us would rather be taxed twenty times at a one percent rate each time than once at 35 percent.

If double taxation is not objectionable as such, what about double non-taxation? Consider in particular the "stateless income" (Kleinbard 2011a) that multinationals are increasingly adept at locating for tax reporting purposes in tax havens where little productive activity occurs. As it happens, the point that zero is not a magic number is already widely accepted. Anti-tax haven rules (both actual and proposed) for resident companies' FSI, in the U.S. and elsewhere, generally do not provide any escape hatch for

charging a very low tax rate that is not quite zero. However, a further issue about global tax minimization through the use of tax havens remains inadequately recognized.

Suppose we observe that Acme Products, in addition to having operations, affiliates, and taxable income in the U.S. and Germany, also has an affiliate in the Cayman Islands that purports to earn a significant percentage of Acme's global profits. We may reasonably suspect that little of Acme's global economic activity is actually taking place in the Caymans, and thus that this is "stateless income," relocated from either the U.S. or Germany by means of tax planning games. Suppose further that, despite having a nominally worldwide system, the U.S. does not get to tax Acme's Caymans income. (We will soon see why, under existing U.S. rules, this may be a realistic assumption.) What should the U.S. think about this phenomenon?

A principle holding that everything should be taxed somewhere once would suggest unconditionally objecting to Acme's profit-shifting to the Caymans. But in evaluating what has happened, there are two distinct possibilities to keep in mind. The first is that profits were shifted out of the U.S., and that Acme has therefore used tax planning to avoid the source-based U.S. tax. The second possibility, however, is that Acme shifted profits out of Germany, thus avoiding the German rather than the U.S. tax. In practice, these two scenarios may be extremely difficult, and perhaps indeed impossible, to tell apart – and one might expect Acme's reported Caymans profits to include some of each.

But should the U.S. regard the two scenarios as equivalent? To answer this question, we must turn to the second big problem with the existing international tax

policy literature, which is its general failure to distinguish between domestic and foreign taxes even though they are very different from the standpoint of a given country.

Second Problem: How Should We Think About Foreign Taxes?

One of the most bizarre aspects of international tax policy analysis is its frequent indifference to the question of whether tax revenues go to the home country that is deciding what international tax rules to apply, or to some other country. As an analytical matter, such indifference is a necessary precondition to focusing exclusively on global rather than domestic tax rates as the key variable. It is at least implicit in the practice of offering foreign tax credits to resident multinational firms. And it is verging on explicit if we do not treat reciprocity as a precondition for offering foreign tax credits. That is, if we allow domestic taxpayers to credit, against their U.S. liability, the taxes levied by countries that do not similarly reimburse their taxpayers for paying U.S. taxes – such as by reason of offering exemption, which (again) is an implicit deductibility system – then we are effectively making a gift of the revenue to foreign Treasuries.

From a unilateral national welfare perspective – that is, one in which we care only about the welfare of domestic individuals rather than everyone in the world, and ignore or assume away other countries' strategic and other responses to our rules – it is clearly erroneous to treat foreign and domestic taxes as interchangeable. After all, the reason we typically regard taxpayers' incentive to reduce or avoid domestic taxes as socially suboptimal is that *we get the money* from those taxes.

Thus, in the purely domestic context, suppose I would decide whether to work for an extra hour based on how much of the earnings I would get to keep after-tax, and that I viewed the tax on these earnings as no less a cost to me than, say, the extra commuting

expenses that I would incur if I accepted the assignment. From the social standpoint, this is an externality problem. I am ignoring the benefit to the domestic individuals who would reap the gain if this extra money went into the public fisc. Analytically, it is the same problem as if I were to ignore the pollution costs that my factory would impose on other individuals by releasing noxious fluids into an adjoining river. With respect to foreign taxes, however, *we don't get the money* – instead, it goes to foreign individuals, whose welfare we commonly disregard in making domestic policy choices (for example, in deciding whether to fund schools at home or abroad).

As noted above, this analysis is incomplete unless we assume, not just that we are exclusively concerned with domestic individuals' welfare, but also that there are no relevant strategic and other interactions between what we do and what other countries do. If, for example, all other countries were resolved to credit our taxes when paid by their own residents if and only if we credited their taxes in the reciprocal setting, then it is possible (depending on further information) that treating domestic and foreign tax payments as effectively equivalent would offer, at least, a plausible rough rule of thumb.

But the fact that there may be important interactions between our foreign tax rules and those in other countries does not establish that we generally should treat domestic and foreign taxes as if they were equivalent after all, from a domestic national welfare standpoint. All it means is that we need to think about those interactions. Treating our treatment of foreign taxes as potentially relevant to how other countries treat our taxes is *not* the same thing as concluding that they are all effectively the same. Consider, for example, that our peer countries have generally shifted towards exemption, and thus towards an effective deductibility system for the U.S. taxes that their resident companies

pay. One can hardly argue that they would abandon foreign tax creditability if we did, when they have already done so without waiting for us to go first.

As we will see, from a purely unilateral national welfare standpoint, mere foreign tax deductibility – including that which results implicitly from having an exemption system for foreign source income – is clearly the right answer, unless and until some further consideration emerges to complicate the analysis. After all, when U.S. people pay foreign taxes, it truly is just an expense from our standpoint, no less than when they pay foreign fuel bills or labor costs, given that we don't get the money.³

There is, however, one last complication that we need to keep in mind. Recall the earlier example where Acme Products, a U.S. firm with domestic and German operations, reported significant profits as having arisen in the Cayman Islands, most likely due to tax planning games that shifted them, for official reporting purposes, out of the U.S. and/or Germany. Despite the ambiguity of from where the profits were shifted, there may be greater reason to suspect that true U.S. income is showing up in the Caymans than that it is showing up in Germany. After all, if U.S. profits are being shifted for official reporting purposes, the taxpayer has more to gain by placing them where the tax rate is actually zero. In addition, as we will see, shifting U.S. profits to a relatively high-tax country such as Germany may simply be the first step, for planning purposes, towards re-shifting them to a tax haven, which often is harder to do directly from the U.S.

The bottom line that this suggests is a bit complicated, and thus needs to be stated carefully. On the one hand, if we regard foreign taxes, unlike domestic taxes, as purely a cost from the social as well as the taxpayer's individual standpoint, then presumably we

should be glad when U.S. companies, owned by U.S. individuals, avoid German taxes by shifting their German profits to the Caymans for German reporting purposes. But on the other hand, when we observe U.S. companies reporting Caymans income, we may have reason to regard this, relative to their reporting German income, as effectively a statistical “tag” that is likely to be correlated with profit-shifting out of the U.S.

There indeed appear to be big numbers attached to this problem. Kimberly Clausing (2011, 1580) estimates that, in 2008, “the income shifting of multinational firms reduced U.S. government corporate tax revenue by about \$90 billion.” This estimate includes a 35 percent gross-up of the overall dollar amount to account for income-shifting by non-U.S. firms that were outside the data set. Such firms are beyond the reach of the U.S. rules for taxing resident firms’ FSI. However, even if one eliminates foreign firms from the revenue estimate, it may suggest that profit-shifting by U.S. firms reduced their 2008 U.S. tax liability by about \$67 billion.

Against this background, if we want U.S. companies to treat foreign taxes as equivalent to any other cost that is incurred abroad, but also want to reduce the incentives for profit-shifting out of the U.S., we face a tradeoff. The former consideration calls for mere foreign tax deductibility, whether the statutory U.S. tax rate for FSI is zero percent, 35 percent, or anything in between. But the latter consideration may call for treating low-taxed FSI less favorably than that which is higher-taxed abroad, even though lowering foreign taxes on true FSI is exactly what we should want U.S. companies, if owned by U.S. individuals, to do.

It is in the nature of the tradeoff that we cannot advance the latter objective without undermining the former objective. After all, we reduce the reward that U.S.

companies reap from overseas profit-shifting if we treat it (in some cases, mistakenly) as evidence of U.S.-to-foreign profit-shifting. Nonetheless, given that this is a tradeoff between competing considerations, it would be a considerable surprise if the unilaterally optimal solution involved having, at any given margin, a company's U.S. taxes increase by a dollar when it reduced its foreign taxes by a full dollar (i.e., the creditability result, whether or not it formally reflected offering foreign tax credits). Instead, the optimal solution might involve companies reaping a worse-than-deductibility, albeit better-than-creditability, marginal result from reporting profits in tax havens.

Third Problem: The Issue of Multiple Margins

In thinking about international tax policy, issues of efficiency are widely agreed to belong at center stage, although distributional issues pertaining to individuals matter as well. But if there is one fundamental rule of clear thinking that one must follow, in order to analyze efficiency in an intellectually coherent fashion, it is that one must proceed, at least initially, by analyzing just *one margin at a time*.

Traditional international tax policy analysis has failed to satisfy this elementary maxim. Instead, a practice of obviously conflating distinct and separable margins has, at least until recently, reigned as unchallenged orthodoxy in the field. Almost everyone seems to agree that the fundamental choice lies between worldwide / foreign tax credit systems, and those that are territorial. This, however, is a compound choice that sloshes together differences at two distinct margins. The first margin concerns *tax rates on foreign source income*. The second concerns *reimbursement rates for foreign taxes paid*.

Tax rates (both average and marginal) on foreign source income – The first, and more obvious, difference between worldwide and territorial systems lies in the tax

rates that they impose on FSI. Under a worldwide approach as commonly conceived, the official statutory tax rate for FSI is the same as that for domestic source income. By contrast, under a territorial system, the tax rate for FSI is zero. Intermediate statutory rates are not so much consciously rejected as ruled out from the start, almost as if they were (for some unknown reason) logically impossible.

As it happens, the scope of the disjuncture may be less than it initially seems, if the worldwide system allows foreign tax credits. Thus, recall the earlier example in which Acme Products' operations and reported income were confined to the U.S. and Germany, with the U.S. taxing FSI at 35 percent but allowing foreign tax credits, while the German rate was 20 percent. If Acme earns \$100 of German income (as determined by both systems) then, after paying \$20 of German tax, it will owe only an additional \$15 of U.S. tax. Thus, its actual average or effective U.S. tax rate on German source income is only 15 percent, as measured on a pre-German tax basis, or 18.75 percent (i.e., $15/80$) as measured after deducting the German taxes.

Both marginal and average U.S. tax rates on FSI may matter for particular purposes. For example, suppose a U.S. company is considering opening a factory either in the U.S. or abroad. The average U.S. tax rate on the factory's expected income may matter more to this decision than the question of how its U.S. taxes would change, in either case, if it earned one additional dollar with all else remaining the same. But once the factory is in place, either here or abroad, the marginal U.S. tax rate may determine how much the company would pay to shift a dollar of profits, for official tax reporting purposes, from inside to outside of the U.S.

Marginal reimbursement rates (MRRs) for foreign taxes – I turn now to the second, analytically distinct difference between the worldwide and territorial systems that traditional analysis fails to consider separately. This is what I call the marginal reimbursement rate (MRR) that they apply to foreign taxes.

In a worldwide system with foreign tax credits, the MRR is in principle 100%. After all, each dollar of creditable foreign taxes that you pay is fully reimbursed by the domestic government via a matching one-dollar reduction in domestic income tax liability. This is what it means to have a 100 percent tax credit.

In practice, two features of real-world foreign tax credit systems may lower the MRR that taxpayers actually enjoy. First, all foreign tax credit systems apply what are called foreign tax credit limits. That is, they provide that, while foreign tax credits can offset the domestic tax liability that would otherwise be due on FSI, they cannot go beyond this point.

To illustrate, suppose Acme Products earns \$10 million in Germany, while having no other FSI. Once again, suppose that the statutory U.S. tax rate for FSI is 35 percent. Absent foreign tax credits, the U.S. tax on Acme's German income would be \$3.5 million. Accordingly, Acme's potentially allowable U.S. foreign tax credits with respect to its FSI are limited to this maximum amount. If Acme paid, say, \$4 million of German tax – as might conceivably happen, even if Germany's statutory rate is lower than ours, if Germany measures German-source income differently than we do – then Acme would end up not being able to claim U.S. foreign tax credits for its last \$500,000 of German taxes. These would instead give rise to a foreign tax credit carryover, potentially allowable in some other year when Acme would otherwise face residual U.S. tax liability

on its FSI. (Since Acme is assumed to have no other FSI, it cannot “cross-credit,” i.e., use the extra German taxes to offset the residual U.S. tax on FSI from a low-tax jurisdiction.)

Permanent disallowance of the extra German foreign tax credits would mean that the U.S. MRR for German taxes switched abruptly from 100 percent to 0 percent when the foreign tax credit limit was reached. However, even if the excess credits became allowable in a future year, the MRR is effectively reduced to something below 100 percent, since Acme suffers a time value loss from the delay of the U.S. reimbursement.

Second, the effective MRR in a worldwide system with foreign tax credits may be lowered by the other really important domestic tax benefit for FSI. This tax benefit is called deferral, as it involves U.S. companies’ getting to defer inclusion of their foreign subsidiaries’ FSI for U.S. tax purposes until it is repatriated, such as through the payment of a dividend to the U.S. parent. This tax benefit is a legacy of realization doctrine, which holds that a change in the value of one’s assets generally has no tax consequences until the gain or loss is “realized,” such as through a sale or other exchange of the asset. Even in the case of a wholly owned foreign subsidiary, which in practice may be little different than a foreign “branch” whose operations the U.S. company conducts directly, realization doctrine causes the profits to be treated as tax-irrelevant to the U.S. parent (which is a separate legal entity) until shareholder-level realization occurs.

In common perception, deferral and the foreign tax credit are respectively what one might call the “Goofus and Gallant” of U.S. international tax policy, referring to a hoary old cartoon series that contrasts bad and good versions of a white middle class suburban, perhaps eight year old, boy. Almost everyone recognizes that deferral, the

“Goofus” of our system, is a terrible rule. In particular, it induces wasteful tax planning behavior by U.S. companies that must jump through hoops to make optimal use of their foreign earnings while avoiding a taxable U.S. repatriation. The reason for retaining it is that, from a pro-territorial perspective, replacing it with immediate U.S. taxation of resident multinationals’ foreign earnings would be even worse. Deferral arose by accident – or, more precisely, through the mindless application of formalistic legal conventions, rather than as a deliberate policy choice – but it has become part of the forced ceasefire-in-place between worldwide taxation and territoriality. I will argue that the foreign tax credit, while it unjustifiably plays Gallant’s noble part in popular perception of our international tax rules, is in fact just as horribly Goofus-like as deferral, leaving aside the analytically distinct fact that each of them happens to lower the effective U.S. tax rate on FSI.

For now, the question of interest is simply how deferral affects the MRR from foreign tax credits. One generally does not get to claim a credit for a given foreign tax liability until the associated FSI is repatriated, and thus becomes subject to U.S. tax. In the interim, however, one’s ultimately allowable foreign tax credits may grow at something like a market interest rate, as additional foreign taxes are paid on the accumulating earnings that remain abroad. This complicates the analysis, relative to the scenario for foreign tax credit limits, which involved straightforward loss of present value for the reimbursement when the claiming of a credit is delayed.

Exactly how deferral thus ends up affecting the MRR for foreign taxes paid is too complicated to discuss in full here, and will be addressed in chapter 3. For now, however, the main bottom line is as follows. One can describe hypothetical conditions

under which the MRR would effectively still be 100 percent, despite the delay in claiming a given credit until the associated income is repatriated. Under realistic conditions, however, the actual expected MRR, and/or that which multinationals' managers will have in mind when they are making decisions, is potentially much lower, and even zero percent in some common circumstances.

In short, deferral can in practice negate the incentive effect that the foreign tax credit would otherwise have (i.e., inducing indifference to foreign tax liabilities). We will see that the reverse holds as well: foreign tax credits can negate the incentive effects that might otherwise result from deferral. And a crucial question in thinking about the future of the U.S. international tax rules is how to think about efficiency of the system as a whole, given the two rules' interactions.

Despite these interactions, however, it is worth keeping in mind the point that, in principle, and except as modified in practice by interactions with other rules, foreign tax credits create a 100 percent MRR for foreign taxes paid by U.S. multinationals. Standing alone, this would result in zero cost-consciousness with respect to foreign taxes. There would be no reason to care whether they were low or high, since they would be subject to full reimbursement by the home government.

A territorial system, by contrast, offers an MRR of zero percent, since it ignores foreign taxes paid. However, of greater note than the fact that the MRR is zero percent is its matching the marginal tax rate (MTR) that such a system applies to FSI. This is why territoriality functions as an implicit deductibility system for foreign taxes. That is, foreign taxes are in effect deductible from pre-foreign tax FSI, and it just so happens that

the remaining amount of such net income is taxed at a zero percent rate, rather than at any positive rate.

Calling territoriality an implicit deductibility system for foreign taxes may initially sound like a cute semantic point. In fact, however, it is an important substantive point. In a deductibility system for foreign taxes (whether implicit or explicit), the taxpayer has an incentive to maximize its after-foreign-tax income, and to treat \$1 of foreign tax liability as equivalent to any other deductible \$1 outlay or forgone \$1 gross receipt. This condition holds whether the tax rate is 0 percent (as under exemption) or any positive number short of 100 percent. Deductibility therefore gives domestic taxpayers exactly the right marginal incentive with respect to foreign tax liabilities, given that the revenues go to people in other countries, rather than to us.

Why the distinction between the two margins matters – Understanding that the choice of domestic tax rate for FSI is analytically distinct from the choice of MRR for foreign taxes really is a key – and perhaps, *the* key – to unlocking the long-frozen terms of international tax policy debate, and deriving fresh insights. A relatively obvious point is that this analytical move encourages considering a much broader range of alternatives in taxing resident companies' FSI. For example, suppose one believes that deductibility offers the correct marginal approach to reimbursing resident companies' foreign tax costs, but that one does not want to raise (or even, perhaps, in the aggregate to change) the domestic tax burden on FSI. One now can see that these two aims are, at least in theory, simultaneously achievable, as one could make foreign taxes merely deductible but simultaneously lower the statutory tax rate on FSI so that the overall burden on outbound

investment, or alternatively overall domestic tax revenues, remained constant (see Clausing and Shaviro 2011).

However, the analytical gain is no less fundamental if one rules out this particular policy response. Even if the worldwide / foreign tax credit and territorial approaches retain an exclusive hammerlock in practice, understanding the two distinct margins that they conflate is intellectually transformative. “Which of these two systems is better?” is simply not a well-posed question, if one fails to grasp its compound character. The questions that one is left with instead, after separating out the two distinct margins – “What should be the domestic tax burden on foreign source income?” and “How should we treat foreign taxes?” – while still not perfect, for example because “foreign source income” is an ill-defined term – are better-suited for coherent analysis.

Fourth Problem: Drowning in “Alphabet Soup”

Another big problem with the existing intellectual debate concerning U.S. (and other countries’) international taxation is that, rather than starting from first principles, such as the welfare economics framework that typically is used in tax policy analysis by economists and their fellow travelers (such as myself), it all too often relies on what I call “alphabet soup” or the “battle of the acronyms.” Under this approach, one starts by describing a single decisional margin at which (all else equal) it would be nice if international taxation were neutral in its effects on taxpayers’ incentives. Each of these competing one-dimensional efficiency norms has its own acronym. During the era of the mainly pro-worldwide consensus, the debate focused on CEN (capital export neutrality), CIN (capital import neutrality), and NN (national neutrality). The recent shift towards a more frequently pro-territorial perspective has featured the addition of CON (capital

ownership neutrality), NON (national ownership neutrality), and GPN (global portfolio neutrality).

Ostensibly, all one needs to do, in order to choose the right international tax policy, is not to consider tradeoffs (as one might have expected from the style of analysis typically deployed elsewhere) but rather to plunge into the alphabet soup, pick a single winner, and determine whether it supports worldwide taxation or exemption. Thus, if one had to say in a single sentence why support for exemption has gained ground recently in academic debate, it would be that, while CEN used to trump CIN, CON has now entered the fray. As we will see, however, a battle-of-the-acronyms approach to international tax policy issues is wholly inadequate.

This single-bullet, alphabet soup approach – if I may be forgiven for mixing my metaphors – would be ill-chosen, and contrary to accepted public economics methodologies outside the field of international taxation, even under the best of circumstances. But the fact that it is used specifically in the international area, which one might have thought somehow related to justifying it, in fact makes using it even worse. Two considerations help to make a battle-of-the-acronyms approach even less suitable in the international realm than it would be as applied to a single country's internal policy choices. I call them the global welfare fallacy and the global convention fallacy.

The global welfare fallacy – A number of the global welfare norms – such as CEN, CIN, and CON – treat global, rather than national, economic welfare as the thing to be maximized. Thus, as we will see, CEN posits that it is a matter of utter indifference whether a dollar of taxes is paid to the U.S. or some other country – even though, from a U.S. standpoint, we don't get the money unless it's our own tax. CON posits that, even if

foreign investors would pay good money to the U.S. for the privilege of getting to invest in U.S. corporations, we should let them do so for free – even though that likewise means throwing away money for no evident reason.

Now, there is nothing fallacious about actually favoring global welfare, rather than purely national welfare – just as there is no fallacy involved if a given individual decides to be a perfect altruist who values everyone else’s welfare equally with her own. Indeed, from a purely ethical standpoint the selfless altruist stands on stronger ground than the rest of us, who fall short of sainthood by caring more about ourselves than about other people, who surely feel pleasure and pain just as strongly, and who matter just as much ethically, as we ourselves do.

But that is not what actually is going on here. Proponents of basing international tax policy on global welfare norms do not in general actually believe that the U.S. should, say, donate three-quarters of its GDP to poorer countries, or afford no preference to its own people in deciding where to build new schools. Instead, when they propound global rather than national welfare arguments, they are confusedly relying on a distinct and much more limited claim, which is that everyone (ourselves included) might end up better off if enough countries cooperated in pursuing global, rather than merely national, welfare.

The underlying point about the possible mutual gains from reciprocity and cooperation, as compared to acting unilaterally and aggressively playing “beggar your neighbor,” is indeed important to international tax policy analysis. However, it calls for more concretely and critically examining how reciprocity and cooperation actually work – which the alphabet soup literature fails to do. Pervasive confusion becomes inevitable

when ostensible global welfare proponents fail to express accurately, and therefore to analyze properly, what they actually care about. For example, one may end up failing to distinguish between areas where mutually beneficial cooperation is more feasible, and those where it is considerably less so.⁴

The global convention fallacy – Suppose we accepted, however, that U.S. international tax policy should be based on maximizing global, rather than U.S., economic welfare. The proponents of alphabet soup often act as if only one further analytical stage is needed: deciding which global welfare norm would be best if everyone adopted it, and then unilaterally implementing it in the U.S.

The problem with this reasoning is that global welfare does not necessarily improve, even if the analysis is otherwise correct, unless other countries generally adopt the same global welfare norm. Thus, as we will see, the U.S. might not actually succeed in advancing the objectives of CEN by pursuing it unilaterally while other countries focused, say, on CIN or CON (even if we were generously to assume that they were thinking in global welfare terms to begin with).

In earlier work (see Shaviro 2007), I posited that the global versus national welfare choice might resemble a prisoner's dilemma in which one is deciding whether to "cooperate" with the other players (by pursuing global rather than national welfare) or to "defect" (via a beggar-your-neighbor version of national welfare), but with the difference that one's choice was partly observable and thus might, to a degree, prompt reciprocal cooperation or defection. From this standpoint, it is perhaps conceivable that the U.S.'s choice along the cooperate-to-defect spectrum would be emulated elsewhere (although this was more plausible when the U.S. was more economically dominant – as it was, say,

in the 1960s – than it is today). But when there are different ways of cooperating, based on distinct global welfare norms, one would need a further assumption, for which there is no contemporary supporting evidence, that other countries will follow the U.S. lead regarding, not just *whether*, but also exactly *how* to cooperate.

Once one recognizes the U.S.'s evident inability to prescribe international tax rules that the rest of the world will follow, alphabet soup analysis rests on implicitly assuming that the question of interest to a given country's policymakers is the following: Suppose all countries were holding a convention to adopt universally binding rules regarding the taxation of FSI. What rule should this convention adopt, from the standpoint of global welfare? No need, apparently, to explain why we should think of the issues this way when no such convention is even remotely imminent.

National welfare norms – But this is entirely enough – or perhaps too much – time spent for now on explaining why we should not bark up the wrong tree by basing international tax policy analysis, intended for a given country, on global welfare norms. Do the national welfare norms that the literature has developed fare any better? As we will see, the answer is no, because they fail to offer any help in addressing the real question of interest, which is how to analyze the tradeoffs between competing efficiency margins, while also including distributional considerations in the analysis.

Fifth Problem: Lack of Conceptual Integration Between Entity-Level Corporate Taxation and the Taxation of Individuals

A final fundamental problem in the field relates to the fact that corporations are the main taxpayers in cross-border enterprise. All significant income tax systems around the world treat corporations as separate taxpayers from their actual owners, who

ultimately are individuals. This is done predominantly for convenience, albeit verging on necessity. Corporations, multinational or otherwise, are not themselves flesh and blood people. They cannot themselves actually bear the burden of paying a tax, nor can they actually live anywhere. However, contemporaneous flow-through taxation of corporate income to underlying individuals, such as the shareholders, is widely considered impractical (see Shaviro 2009a, 154).

For this reason, worldwide residence-based corporate taxation unavoidably depends on corporations', rather than individuals', deemed residences. This proves to be of central importance in thinking about international taxation. So does the fact that legally separate corporate entities – say, a U.S. parent and its wholly owned foreign subsidiaries – are treated in many respects as separate persons, even though they are not in truth distinct like two human beings.

Yet corporations were long the key missing player in international tax policy analysis, which largely proceeded as if it made no difference that, for example, the “Americans” potentially being subjected to U.S. worldwide taxation were companies that had filed incorporation papers here, not people who were citizens or residents. More recently, as we will see with regard to the more recently coined acronyms CON and NON (which focus on the importance of which company owns a particular asset), this has begun to change, and corporations that are taxed as separate entities have at last taken their indispensable central place on the analytical stage.

While corporations are the principal direct actors and taxpayers in the international realm, they still are not (and cannot be) the true ultimate parties of interest. Thus, important and laudable though it is that the international tax policy field has begun

in recent years to focus on the significance of their being the main taxpayers in the international realm, there has not as yet been much progress on the next step, which involves relating entity-level taxation back to the individual level. Here the question of interest is how international taxation relates to the distribution issues in tax policy.

Thus, suppose the U.S. is deciding how to tax U.S. multinationals. In so doing, it indirectly taxes the owners, all of whom ultimately (once one looks through all of the intermediate legal entities that own financial interests in the firms) must be either U.S. or foreign individuals. Accordingly, in evaluating how we tax U.S. multinationals, one must evaluate the possible effects on both of these two groups. This stage has generally been missing in analysis to date, and I plan to show its relevance in this book.

Towards a Better Analysis of International Tax Policy Issues

Having spent so much time summarizing what I believe is wrong with the existing international tax policy literature, it is only fair that I acknowledge the old schoolyard dare, “put up or shut up.” In other words, how do I propose to improve the analysis? In general, I believe that this mainly involves addressing the following five, to a degree overlapping, questions.

1) What companies should we define as U.S. residents? – When we say an individual is a U.S. resident, we are defining her as a member of our community, a status that may be associated with our conveying both benefits and burdens. By contrast, when we define a corporation as a domestic resident for purposes of our international tax rules, the main thing it means is that we can, if we like, tax what we call its FSI.

As we will see, we would have good reason to tax the entire worldwide income of every company in the world – wholly without regard to whether it had connections of any

sort to the U.S. – if this were legally, politically, and administratively feasible, as of course it is not. In particular, this approach would have the appealing feature, from an efficiency standpoint, of making the source-based domestic tax wholly unavoidable through locational choices (such as by investing abroad rather than here).

The question presented, then, is what subset of all the companies in the world, as among those that we *can* indeed define, if we choose, as U.S. residents, we should decide to make potentially subject to U.S. worldwide taxation. To what extent are the benefits of taxing all companies on all of their worldwide income merely scaled down, rather than lost, when we can only tax some of them? In choosing which attributes to use in identifying resident companies, a key question is tax-elasticity. Thus, for clearly feasible attributes, such as one's place of incorporation or headquarters location, how flexible are taxpayers prospectively, and how easily can they alter what they have already done? Moreover, what attributes might we be reluctant to tax-discourage, such as on the ground that their selection by corporate actors might have positive externalities for people in the U.S.?

2) How should we define foreign (and domestic) source income? – In a world where governments claim exclusive authority within geographically limited realms, a source concept for income taxation verges on being inevitable. A country can be expected to at least consider taxing foreigners' "inbound" income, given that it has the power to do so. And for one's own residents, even if one imposes worldwide taxation, the question of whether a given dollar was earned at home or abroad is potentially of interest, given that one's control (and that of other governments) differs as between the two realms.

Unfortunately, the source of income is not a well-defined economic idea (Ault and Bradford 1990). Perhaps its meaning is clearest when a given individual, who is in a particular place, creates economic value by doing things exclusively there. Even if she has foreign sales, the work of production took place entirely at home. But even this simple case rapidly runs into several conundrums. For example, suppose people who are living abroad helped her to generate the foreign sales. Or suppose that the overseas “market” countries want to claim some of her income from selling items to their residents. What is more, even just for individuals, the inquiry into the source of income rapidly gets harder once we start thinking, for example, about where the income from financial instruments was earned, or about how to treat expenses that arguably relate to gross income earned in multiple locations. When we shift from the world of individuals to that of multinational corporate entities, the source inquiry gets exponentially harder still, and less resolvable even in theory.

So far, I have been treating source determinations as purely a factual inquiry into some underlying actual or partial truth. To the extent that there are correct answers about source, implementing them through one’s international tax rules serves economic efficiency, from a national welfare perspective, at one particular margin. It avoids either encouraging or discouraging domestic taxpayers from integrating domestic operations with those conducted abroad. By contrast, if one mismeasures the domestic income component of cross-border operations, one may inefficiently encourage or discourage cross-border operations just for tax reasons.

Unfortunately, however, accuracy-based source determinations face the problem that there often is no there there. In addition, accuracy may not be the only value that

matters from the standpoint of a given country's self-interest in setting international tax rules. Two further potentially important considerations are (a) deliberately discriminating between types of income, and (b) affecting what I call taxpayers' sourcing electivity (their ability to engage in cheap or even cost-free income-shifting for tax purposes).

Discriminating between types of income – As we will see, countries may have good reason to use their source (and related) rules as a device for treating relatively mobile income more favorably than that which is more geographically fixed. The greater an item's geographical mobility, the worse the tradeoff between revenue-raising potential and the efficiency cost to the taxing jurisdiction of attempting to reach it. In principle, one could respond to this consideration by expressly applying lower tax rates to mobile than immobile income (albeit, taking account as well of the possibility that taxpayers may shift between the two types). However, responding accommodatingly to taxpayers' use of tax planning tricks to shift profits purely for official reporting purposes, when these tricks pertain to economically mobile income, may accomplish the same end with greater political and administrative convenience.

The principle that one may benefit from discriminating between types of income also applies to that which one classifies as FSI. For example, income that is reported as arising in a tax haven, or that is highly mobile for reporting purposes and thus appears likely to be shifted to tax havens, often faces current taxation, in the hands of resident multinationals, even if one does not directly challenge their classifying it as FSI. This reflects that, from a practical standpoint, the source determination that one accepts for such income may matter less than how one actually taxes it.

Cheaper versus costlier sourcing electivity – Now suppose we are focusing on one particular type of income, rather than on how its level of taxation compares to that for other income. A country’s international tax rules can make it easier or harder – cheaper or costlier – for taxpayers to succeed in treating it as foreign source rather than as domestic source. An important criticism of the existing U.S. source rules, such as in Edward Kleinbard’s (2011a) discussion of the “stateless income” phenomenon, is that they create too much electivity, by making legally successful profit-shifting too cheap rather than too costly. This amounts to arguing as well that there is currently too much discrimination (even if the optimal amount is not zero) in favor of income that is earned by multinational companies.

In this book, while not attempting to provide definite suggestions for U.S. rule design, I will show that the cost-of-electivity framework can play an important role. Two particular issues stand out. The first is how to deal with debt (and thus the interest deductions) incurred by members of an affiliated worldwide group of companies. Economically speaking, the details both of intra-group debt and of who within an affiliated group is the borrower with respect to third party debt, may matter very little, yet the tax implications may be great. The existing rules thus create very substantial sourcing electivity that probably ought to be reduced.

The second big electivity issue concerns the tax treatment of intra-group transactions, such as sales. Suppose, for example, that a U.S. member of a multinational group either sells productive inputs to a foreign affiliate in connection with an export, or buys such inputs with respect to an input. The deemed intra-group sale price (known as

the transfer price) plays a key role in determining U.S. tax liability, despite the fact that it may not matter economically at all (since the group members are commonly owned).

Transfer pricing rules typically rely on the notion of a “comparable arm’s length price” between unrelated affiliates. As has been widely discussed, however (see, e.g., Durst 2007), this concept is bad enough in theory, and even worse in current practice. Many have therefore proposed replacing transfer pricing with the use of an approach called formulary apportionment, under which the affiliated group’s worldwide income would be apportioned through the use of a mechanical formula, typically based on the relative proportions of its sales, payroll, and/or tangible property in each country. This approach, however, is widely criticized as well, based both on the continued distortions and inaccuracy that it would entail and on the political difficulty of devising a common formula that all nations might agree to apply. Once again, however, a focus on the issue of cheaper versus costlier electivity can at least help to clarify the tradeoffs.

3) Why should we want to tax (or not tax) resident companies’ foreign source income? – Suppose that we have now defined both domestic source and foreign source income, along with substantively related questions such as that of which member of an affiliated group earned a given dollar of income. Taking all this as given – although in fact, our preferred definitions of domestic source income and FSI might partly depend on how each is treated⁵ – what motivations would one have for imposing a positive tax rate on resident companies’ FSI?

The core reason should be obvious. If we also take as given the imposition of a source-based tax on domestic income (for reasons that I discuss in chapter 5), one faces a fundamental incentive problem. Absent a tax on FSI, all taxpayers can avoid the source-

based tax by investing or reporting income abroad, rather than at home. But for anyone whose FSI is also taxed, and indeed under the same terms, the source-based tax becomes unavoidable through locational choices.

This is why, as noted above, at this margin it would clearly benefit the U.S. (or any other country) to tax all companies in the world on all of the income that they earned anywhere in the world. But this of course is impossible and would likely provoke adverse responses from other countries. So the potential benefit reduces to that available from taxing just resident companies on all of their worldwide income.

Further motivation for taxing FSI relates to the shareholder level. The U.S. imposes a worldwide income tax on individuals. The worldwide character of this tax clearly is necessary if we are using income as a metric for some underlying distributional concept such as “ability to pay.” For example, one cannot really tell how well-off Bill Gates or Mark Zuckerberg is, compared to people who are less materially fortunate, unless one counts all of their income, whether earned at home or abroad.

But with corporate income being taxed only at the entity level until there is a shareholder-level realization (such as from receiving dividends or selling stock), the tax on FSI becomes too easy to avoid unless we also tax the corporations in which U.S. individuals own stock on a current and worldwide basis. To be sure, here we face the same problem as that arising under the efficiency heading of making the source-based tax unavoidable. That is, the logic calls for taxing *all* companies’ worldwide income (insofar as they have U.S. shareholders), but we can only impose it on companies that we define as U.S. residents.

A further potential motivation for taxing FSI relates to foreign shareholders. Again, suppose that we could tax them on their worldwide income no matter where earned, and no matter how (such as through a foreign company) they invested. This would potentially lead to wealth transfers from them to us. Even if we can only do this when they invest through U.S. resident companies, the possibility of such a wealth transfer remains. The key question, as we will see in chapter 5, is whether they are willing, in effect, to pay (through imposition of the U.S. worldwide tax) for the privilege of investing through a U.S. company.

If these are the motivations for taxing U.S. companies' FSI, what are the motivations for not doing so? To explain, suppose initially that U.S. corporate residence status, no matter how we defined it, would be entirely elective for all taxpayers. That is, it would have no practical effect on them of any kind, except for its causing their non-U.S. income to face a U.S. tax. Under these circumstances, the tax on resident companies' worldwide income would be expected to raise zero revenue. It would be pointless – albeit not affirmatively harmful, since no-one would either elect it or be adversely affected by opting out.

Now suppose instead that U.S. corporate resident status is not entirely elective. That is, some people would be willing to pay for it in the form of expected higher U.S. taxes. Now the revenue estimate is no longer zero, but there are also potential efficiency costs to think about. In particular, U.S. individuals may incur deadweight loss to avoid it. And if U.S. people enjoy any positive externalities when investors choose the attributes that cause a company to be classified as a U.S. corporate resident, these may be reduced

by the tax disincentive to choose them. All this, however, merely leads us to the next question, which is what tax rate we should apply to U.S. companies' FSI.

4) How should we determine what average / effective tax rate to apply to resident companies' foreign source income? – As noted above, the existing international tax policy literature generally contemplates just two possible domestic statutory tax rates for FSI: zero percent or the full domestic rate. However, even strongly committed proponents of worldwide U.S. taxation, without deferral but with the allowance of foreign tax credits, evidently agree that the average or effective U.S. tax rate on FSI should be somewhere between those two endpoints.

How should we think about the criteria for determining optimal U.S. tax rates for U.S. companies' FSI? Fortunately, the existing tax policy literature – although not its expressly international component – can help in addressing this question. A starting point is the optimal commodity tax literature, founded by Frank Ramsey (1927). This literature addresses efficient rate-setting if there are excise taxes on various consumer goods. It finds that varying, rather than uniform, rates are generally optimal when the potentially taxable items differ in their tax-elasticity. More particularly, the “inverse elasticity” rule holds that, the more tax-elastic a given item, the lower the tax rate on it should be, all else equal. Further and related insights can be gleaned from the existing tax policy literature concerning the marginal efficiency cost of funds (Slemrod and Yitzhaki 1996). Here a key idea is equalizing the ratio of deadweight loss to revenue raised as between alternative instruments (such as the tax on domestic source income as compared to FSI), in light of a common tendency for this ratio to grow worse, for any given instrument, as one uses it more (such as by raising the rate).⁶

To illustrate optimal commodity tax analysis (and how it relates to the marginal efficiency cost of funds), suppose a country needs to raise a fixed amount of revenue, with as little economic distortion as possible, solely through commodity taxes on two items – say, milk and orange juice. Suppose that these are the only available drinks apart from tap water, which let's assume cannot be taxed.

A narrowly focused neutrality-based approach might call for taxing both milk and orange juice at the same rate, so that people do not switch between milk and orange juice for tax reasons. However, while this consideration is indeed relevant to the choice of tax rates for the two items, it may be trumped by other considerations. In particular, suppose that orange juice consumption is much more tax-elastic than milk consumption, perhaps because people are more willing to switch from orange juice to tap water. Then it turns out, as the optimal commodity tax literature shows, that the tax rate on milk should be higher than that on orange juice. And what should be equalized, rather than the two items' tax rates, is the marginal efficiency cost, in terms of deadweight loss, of the last dollar of revenue raised from each of the two levies.

In chapter 5, we will see that this mode of analysis is indeed relevant to international tax policy, with domestic source income taking the place of milk while resident companies' FSI takes the place of orange juice. (Investing abroad through foreign companies takes the place of drinking tap water.) This is why overall optimization – even if defined in terms of efficiency concerns *plus* those of distribution – depends on something very different from picking a particular neutrality (such as that either between orange juice and milk, or orange juice and tap water) and basing one's

whole policy on it. Rather, it requires rejecting any analysis that is based on neutrality just at one particular margin.

Unfortunately, the optimal commodity tax and related literatures are unlikely in practice to yield definitive answers concerning the optimal U.S. tax rate on FSI. Yet they not only help clarify the terms of inquiry, but also may support general answers within broad ranges. For example, they help to show why exemption's zero tax rate is unlikely to be optimal. If one started with tax rates of zero for FSI and some positive rate for domestic source business income, it is highly likely that gradually raising the former rate and lowering the latter rate, in a manner that was revenue-neutral overall, would have at least initial efficiency benefits, even if these disappeared well before the point at which the two rates had been equalized.

5) How, or in what manner, should we tax resident companies' foreign source income? – A standard maxim in tax policy outside the international arena holds that it generally is optimal to broaden the base and lower the rates. Taking a particular revenue target as given, this approach typically yields two distinct efficiency benefits. First, base-broadening reduces the distortions that would result from treating some items more favorably than other items (assuming, of course, inadequate justification for the disparate treatment, such as from externalities or elasticity differences). Second, lowering the marginal rate may in particular circumstances reduce tax distortions, such as from discouraging work and saving, even if the average rate remains the same given base-broadening. (In other instances, however, only the average rate matters.)

Applying this maxim to the taxation of FSI would suggest the likelihood of efficiency gain from a shift in policy that involved (1) eliminating deferral and repealing

the foreign tax credit, plus (2) lowering the tax rate on FSI so that some underlying target, such as revenue raised or the average tax rate for FSI, remained the same. This point never seems to get mentioned, however (other than in my prior work), probably for three main reasons. First, the straitjacket of conventional thinking seems to rule out any statutory tax rates for FSI in between zero (or something close to it) and the full domestic rate. Second, pervasive confusion regarding how to think about foreign taxes impedes thinking about the foreign tax credit as an inefficient tax preference – as it clearly is, from a unilateral domestic standpoint. Third, the shift would violate dozens of bilateral tax treaties that the U.S. has signed, and might also raise additional administrative and political economy concerns.

Even if this last objection is dispositive, however, it is worth knowing whether the general base-broadening principle would otherwise apply here. For example, it might affect how we think about the range of choices that are indeed realistically available. (Territoriality is a treaty-compatible approach to repealing deferral and the foreign tax credit that merely happens to feature a lower tax rate for FSI than that which would otherwise be optimal.) But rather than simply assuming the applicability here of the general base-broadening principle, one should more particularly examine how deferral and the foreign tax credit affect taxpayers' incentives and behavior in practice.

With regard to foreign tax credits, we have already seen that a 100 percent MRR would cause U.S. taxpayers to be wholly non-cost-conscious with respect to foreign taxes paid, likely to the detriment U.S. national welfare given the lack of general credit-granting reciprocity by other countries. And while often the MRR may be significantly lower than 100 percent in practice, mainly due to deferral, inefficiency still results from

treating foreign taxes more favorably than other foreign outlays (although the “tagging” point about low foreign taxes must be kept in mind as well). Accordingly, the foreign tax credit is an inefficient tax preference that ought to be eliminated, holding to one side the question of what should be the overall tax burden on FSI.

What about deferral? As we will see in chapter 3, there are hypothetical conditions in which it would not actually have any distorting effect on U.S. multinationals’ decisions regarding when and whether to repatriate funds that (tax considerations aside) could better be used at home. Unfortunately, however, these conditions do not apply in practice, and thus deferral causes large distortions that appear to be getting worse over time. The maxim that one should broaden the base and lower the rate therefore applies here as well.

Finally, what about a system that has both deferral and foreign tax credits? As noted above, each tends to blunt the other’s worst incentive effects. However, a system featuring both of them is guaranteed to impose extremely high tax planning and compliance costs relative to the revenue raised. An awful result would not, however, be similarly inevitable if the system instead applied an expressly lower tax rate to FSI. Even the fact that source determinations would now explicitly be necessary in all cases, rather than just when foreign tax credit limits potentially apply, would matter very little in practical terms. After all, under present law, U.S. companies pervasively must make determinations that have substantively similar effects, in the course of determining whether a given dollar of income belongs to the U.S. parent or to a foreign subsidiary.

So long as the U.S. international tax rules makes use of foreign tax credits and/or deferral, efforts to design or analyze reform proposals are confounded by what I call the

“iron box.” Cut back on foreign tax credits and you make deferral worse; cut back on deferral and you make foreign tax credits worse. A step forward as to either of these inefficient rules triggers a step back as to the other of them, even if the latter’s technical terms have not been changed. Plus, one raises the separate question of whether the overall tax burden on FSI is being moved closer to, or further from, its optimal level. So long as U.S. international tax policy remains trapped in the iron box, significant overall progress is hard to achieve, and even marginal improvement is hard to judge.

In illustration, suppose we repealed deferral and went to a pure worldwide system with foreign tax credits. Even apart from the question of whether (absent other changes) this would unduly increase the effective U.S. tax rate on resident companies’ FSI, this would result in worsening the incentive effects of the foreign tax credit rules. On the face of things, it might look as if the those rules were unchanged. But in fact, U.S. companies would now have an effective MRR of 100 percent in *all* cases other than those where the foreign tax credit limit might apply. They would thus have a greatly reduced incentive, relative to today’s law, to seek to reduce their foreign tax liabilities.

Now suppose instead that we changed the foreign tax credit in such a way as to make it much harder for U.S. companies to engage in cross-crediting (i.e., the elimination of any U.S. tax on repatriations of low-taxed FSI by pairing them with the repatriation of high-taxed FSI).⁷ While this might effectively reduce U.S. companies’ MRR with respect to foreign taxes, the resulting U.S. national welfare benefit would come at the cost of increasing the distortions associated with deferral, by making tax-free repatriations harder to arrange (Sullivan 2009).

Only by eliminating *both* foreign tax credits and deferral can we break out of the iron box that currently constrains U.S. international tax policy. And again, this need not imply increasing the effective U.S. tax rate on resident companies' FSI – the desirability of which depends on wholly separate considerations.

If We Now Know the Right Questions, What is the Answer?

So much for the windup – what, then, is the pitch? Under the analysis that I am suggesting, what should the U.S. actually do in the international tax policy arena?

The answer to this question is complicated by the fact that it depends on how we interpret the choice set. How constrained are we in picking the best available solution? Consider, for example, the two core reasons why our international tax rules are so horrendously bad at present. The first is that we practice entity-level corporate income taxation, while the second is that existing worldwide systems use deferral and foreign tax credits. One cannot make a proper recommendation without knowing to what extent either or both constraints are binding.

As we will see in chapter 6, it is conceivable that the core dilemmas in current U.S. international tax policy could actually be eliminated. For example, replacing the income tax with a progressive consumption tax, depending on its form, could make both corporate residence determinations and transfer pricing wholly unnecessary. And even if the income tax remains in place for individuals, Alan Auerbach's (2010) recent proposal to replace the corporate income tax with a particular type of consumption tax could have the same effect. Even this, however, arguably is a more decisive Gordian knot solution to the existing dilemmas than we should realistically expect.

Now suppose that we are indefinitely stuck with entity-level corporate income taxation. What should we do then? Here, while we are deep in the realm of unappetizing and uncertain tradeoffs, I will end up (in chapter 6) suggesting the adoption of a modified territorial system, even though a zero tax rate for U.S. companies' FSI is likely to be too low. The main reason would be to eliminate foreign tax credits and deferral in a politically feasible and treaty-compatible manner, perhaps without requiring the effective tax rate on FSI to drop quite as low as one might initially have assumed was necessary under a system that is called "exemption."

It appears to be widely recognized that treaty compatibility does not actually require a zero rate on resident companies' FSI. Suppose – although the point is legitimately debatable – that 5 percent, or perhaps even 8 or 10 percent, is acceptable under an "exemption" system, if rationalized as an indirect way of taxing domestic source income that has been mislabeled. Then the horns of the dilemma become considerably less sharp, because one need neither violate treaties nor pre-commit to under-taxing FSI in the course of breaking out of the iron box.

In addition, as is common in putatively territorial systems, resident companies' foreign profits that they treat as arising in tax havens can be taxed at a higher rate than other FSI, albeit subject to the offsetting concern about undermining the companies' foreign tax cost-consciousness. Quasi-exemption's repeal of deferral might also be accompanied by imposing a transition tax on U.S. companies' pre-enactment unrepatriated foreign earnings, which they accumulated under the expectation that they would face a U.S. tax at some future point.

Further associated changes that my analysis will support (or at least show to be highly plausible) include expanding the U.S. corporate residence rules, and improving the source rules. In this regard, one important step, pertaining especially to debt, would be to reduce as much as possible the legal significance of formal entity lines within an affiliated corporate group. I will also argue that there may also be decent grounds to move part-way towards a recent proposal (Avi-Yonah, Clausing, and Durst 2008) to replace transfer pricing with a version of formulary apportionment that gives the sales factor greater weight than property and payroll.

The rest of the book will more fully explore and develop the ideas set forth in this chapter, but with a very different organizational scheme. Whereas here I proceeded straight from the core dilemmas in international tax policy, to the defects in prevailing modes of analysis, to my own proposed analysis and its main policy implications, in the rest of the book I begin by taking a step back. In-depth discussion requires a thorough grounding in the general details of the current U.S. international tax system. Thus, Part 1, consisting of chapters 2 and 3, first reviews the basic U.S. international tax rules, and then addresses in greater detail the design challenges that they raise, along with their main incentive effects and planning implications. Part 1 could be skimmed or even skipped by readers who either are already well-versed in the operational details, or else do not wish to delve too deeply into the U.S. international tax system's plumbing. However, it does (chapter 3 in particular) develop some points that are important to the subsequent analysis.

Part 2 then shifts to a broader policy focus. To this end, chapter 4 addresses the global welfare perspective on U.S. international tax policy. Chapter 5 addresses the

unilateral national welfare perspective. Finally, chapter 6 addresses the question of what practical steps might be taken to improve U.S. international tax policy.

ENDNOTES

¹ During the 2011 debt ceiling fight between the Obama Administration and Congressional Republicans, Treasury Secretary Geithner reportedly stated that “[t]he goal is territorial” and that the Obama Administration was “prepared to move off decades of Democratic orthodoxies” by agreeing to go “95 or 96 percent” of the way there (Woodward 2012, 242). During the 2012 presidential campaign, however, the Obama Administration endorsed research suggesting that shifting to a territorial system, as proposed by Governor Romney, would cost 800,000 U.S. jobs (see Clausing 2012).

² I first laid out pieces of this critique in articles such as Shaviro (2010) and Shaviro (2011b). While my critique of the existing literature may sound harsh, I am not alone in making it. Thus, consider the title of Graetz (2000), an article that stands as an important precursor to my work here: “Taxing International Income: Inadequate Principles, Outdated Concepts, and Unsatisfactory Policies.” Or consider the statement in Auerbach (2010) that, while worldwide and territorial systems “are generally presented as the two reform options available the United States could benefit by moving beyond these two heavily discussed choices.”

³ It is true that the U.S. generally benefits if peer countries are able to maintain healthy and well-functioning tax systems. Accordingly, rampant tax avoidance that undermined their fiscal capacity might also harm us. However, this is a collective action problem – since all peer countries may benefit from each others’ general fiscal health – that calls more for considering multilateral cooperation than for our unilaterally subsidizing tax payments to peer countries by resident U.S. companies. We may also benefit if peer countries maintain healthy economies with close to full employment and good wages.

Yet it seems unlikely that we would want to respond by making U.S. companies' foreign wage payments better than deductible.

⁴ For example, I will argue in chapter 4 that multilateral cooperation is considerably more feasible with regard to information-sharing and perhaps the “stateless income” or tax haven problem than with regard to reimbursing source-based taxes through the allowance of foreign tax credits.

⁵ For example, the greater the tax benefit from re-labeling domestic source income as FSI, the tougher the source rules might need to be.

⁶ As I discuss in chapter 5, the optimal tariff literature also provides important and related insights. This literature shows that a country can benefit from levying tariffs on imports when it has significant market power – in particular, the ability to affect the world price for a particular commodity. Chapter 5 will show that issues of market power, like that which the optimal tariff literature emphasizes, are more generally central to the question of how heavily (if at all) both domestic source income and resident companies' FSI should be taxed.

⁷ As we will see in chapter 3, the Obama Administration in 2009 proposed a so-called “pooling” rule for foreign tax credits that likely would have this effect.