

A PROPERTY RIGHTS APPROACH TO TEMPORARY WORK VISAS

Alessandra Casella and Adam B. Cox*

Temporary labor visa rules in the United States are criticized on three grounds: for failing to allocate visas efficiently, for failing to adequately protect domestic workers, and for exposing migrant workers to exploitation. We argue that it is possible to address all three problems by re-configuring the property rights associated with the visas and carefully designing the mechanism for allocating those rights. Our core insight is to unbundle the two rights that today are typically combined: the firm's right to employ a foreign worker, and the worker's right to reside and work in the country during that time. A three-pronged approach— auctioning abstract pre-contract visas to firms, allowing their trade on a secondary market, and transferring the visa's ownership to the worker upon signature of the employment contract—has the potential to improve the efficiency of visa allocation, to better shield domestic workers, and to protect foreign workers from exploitation.

* Professor of Economics and of Political Science, Columbia University; Robert A. Kindler Professor of Law, New York University School of Law. We are grateful for comments and suggestions from Ryan Bubb, Scott Hemphill, Daryl Levinson, Eric Posner, Stefan Schlegel, an anonymous editor, and participants in the NYU Faculty Workshop, as well as the Guest Worker Conference at the University of Chicago Law School.

INTRODUCTION

Visas are valuable property rights. Yet this basic insight is too often overlooked in debates about the structure of America’s labor migration rules. This is particularly true for temporary work visas, which are interesting (and importantly different from green cards) because they actually bundle together two separate legal rights. The first is the right for a firm to employ a foreign worker in the United States for a specified period of time. The second is the worker’s right to reside and work in the country during that time.

The two rights are combined because American immigration law delegates to employers much of the power to screen labor migrants. Employers select overseas workers and petition for their entry. And if the petition is approved and a visa is granted, that visa fuses together the rights of the firm and the worker.

The manner in which temporary visa rights are bundled and distributed has contributed to three persistent complaints. First, these valuable rights are allocated to firms inefficiently. Many programs operate on a first-come, first-served basis. And when programs are popular—as is the case for the well-known H-1B program for highly-skilled workers—visas are allocated essentially by a lottery. Second, the allocation and screening rules do little to protect the domestic workers they are purportedly designed to shield. By giving away the right to employ overseas workers for free (or effectively for free), the code encourages employers to use temporary visas to secure cheaper labor. Arcane “labor certification” rules formally prohibit such behavior, but are widely viewed as toothless in practice—and perhaps even unenforceable in principle. Third, the structure of temporary visas exposes migrant workers to exploitation. Because temporary labor visas typically tie an immigrant’s right to reside in the country to her continued employment by the sponsoring employer, that employer has tremendous leverage. If the worker quits or is fired, she immediately falls out of status and is legally obligated to depart the United States.

In this article, we argue that it is possible to address all three of these problems by re-configuring the property rights associated with a temporary labor visa and carefully designing the mechanism for allocating those rights. Our core insight is to unbundle the two rights that today are bundled, treating these rights differently before and after an employment contract is signed—that is, thinking of temporary visas as having a different nature before and after the employment relationship is entered into. Our proposal has two signal virtues. First, it allows us to exploit market forces to improve the allocation of visas, without converting workers into commodities or permitting immigrants themselves to purchase the valuable right to reside in the United States. Second, it draws on regulatory tools that are routinely used and, in some cases, already nascent in existing immigration law. Thus, our proposal could be adopted without requiring radical alterations to America’s existing rules.

I. TEMPORARY LABOR VISAS IN THE UNITED STATES

Temporary visas are central to labor migration policy in America today. This might come as a surprise to anyone who casually perused the Immigration and Nationality Act. The Act is organized around the process of permanent migration, with temporary migration treated formally as an exception to the typical migration path. But most labor migrants do not enter the United States as lawful permanent residents. They come as temporary workers.

There are three reasons for this. First, relatively few green cards are handed out each year on the basis of labor market qualifications. The vast majority of green cards are allocated, instead, on the basis of family relationships. Second, today temporary labor visas admit large numbers of workers to the country each year. As Table 1 shows, in 2014 more than 300,000 visas were handed out to noncitizens under two of the most numerically significant temporary labor visa categories—H-1B visas reserved for high skill workers, H-2 visas reserved for low-skilled workers. In contrast, only about 150,000 immigrants received green cards that year under *all* categories of permanent labor migration combined. Third, most immigrants who are admitted to permanent residence under one of the employment-based visa categories *are already living* and working in the United States—typically here on temporary labor visas.

[Insert Table 1 here]

The regulation of temporary work visas is the result of long-standing political compromises. First, only certain types of workers are admitted. That is part of what produces the proliferation of temporary labor visa categories in Table 1. And as that Table illustrates, current rules for admitting temporary workers favor high-skilled over low-skilled labor. Second, American immigration law restricts the number of workers who can come each year. Many temporary admissions categories come with annual quotas, so an otherwise-eligible noncitizen will receive a visa only if the quota for her category of admission has not yet been reached.¹ Third, immigration law restricts workers' terms of admission, imposing various restrictions on the worker's stay in the country. A visa for seasonal agricultural workers, for example, permits admission only for the purpose of doing "seasonal work," provides for a maximum term of admission of one year, and prohibits the worker from changing employers after entering the country. The H-1B visa for high-skilled workers, on the other hand, provides for longer terms of stay and somewhat greater job flexibility.²

The restrictions on temporary labor admissions are supposed to serve three goals: first, providing access to overseas workers for sectors where those workers are most needed; second, protecting domestic workers from the adverse effects of increasing the supply of labor in the United States; and third, preventing the exploitation of workers employed in the

¹ The details of the quota system are actually far more complex: in some cases (but not all) unused visas from one category will be allocated in a future year to the same category or to some other category; a separate quota restricts the number of visas that can be issued in any given year to nationals of a particular country; and so on. But these complications apply principally to the categories of permanent migration, so they are irrelevant for our analysis.

² An H-1B visa holder can begin working for a new employer once that employer files an H-1B petition for that employee (8 U.S.C. § 1184(n)). But if an H-1B visa holder is fired by her current employer, she immediately falls out of status and is legally obligated to leave the country. The same is true if she voluntarily quits before a new employer has filed an H-1B visa petition on her behalf. Moreover, if an H-1B worker transfers from an quota-exempt employer, such as a university, to an non-exempt employer, the new employer's petition counts toward the annual visa cap. Because in recent years the quota has been reached within days of April 1, this rule significantly limits the ability of H-1B employees working for exempt employers to take jobs with non-exempt employers.

United States on temporary visas. Unfortunately, the current rules are often thought to do a terrible job on all three fronts.

A. The Inefficiency (and Unfairness) of Visa Allocation

Under nearly all of the temporary labor visa categories, employers, not the immigrants themselves, petition for visas for prospective employees. And visas are handed out upon demand, without regard to employer need or social welfare, until the annual quota for the visa category is exhausted (INA § 214(g)(3)).

For any visa category where demand regularly outstrips supply, this leads to a race among employers, and to the unraveling of the allocation scheme. Under the H-1B program in 2016, for example, a record 236,000 applications were received during the first week that employers were permitted to submit applications, vastly exceeding the relevant quota of 85,000 (O'Brien 2016). When this happens, as it has in every recent year, the government selects applications by lottery from among the pool of all valid applications filed during a one-week filing period that begins on April 1. No cap-eligible applications are accepted after that point. In 2016, that meant that each application had a roughly one in three chance of success.

The lottery creates incentives for firms to engage in wasteful strategic behavior. In the H-1B system, for example, there are widespread reports that firms submit multiple petitions for what is effectively the same job, in order to increase the odds of having one of their petitions drawn in the lottery. Moreover, the current rules permit a single employee to appear as prospective hire on applications from multiple firms, which enables large firms to have more than one subsidiary petition for the same employee.

Commentators complain that, in addition to being inefficient, the current system has undesirable distributive consequences. The complexity of the H-1B process favors firms that are larger, firms that have prior playing experience, firms that are more legally sophisticated, and firms that intend to hire larger numbers of H-1B visas. In 2014 for example, more than 10,000 companies applied for the available 85,000 H-1B visas, but more than a third were awarded to just 20 companies (Park 2015). Thus, the system currently gives away a valuable, government-provided labor market benefit—essentially for free—to a surprisingly small number of firms.

A large proportion of these successful firms deploy a business model that exacerbates the pressure of immigration on competing domestic workers. So-called “outsourcing” firms received a large fraction of H-1B visas last year, with 20 percent of all available visas going to outsourcing companies headquartered in India, and more than 5,500 going to a single firm, Tata Consultancy Services (Park 2015). Tata and other outsourcing firms specialize in providing temporary foreign workers, typically on very short term consultant contracts. The contention is that these workers spend their short time in the United States learning the skills of a particular job currently done by a domestic worker—often being trained by the domestic worker himself. The job then leaves the United States with the overseas worker, being outsourced to the foreign office of the outsourcing firm that sponsored the foreign worker’s entry into the United States.

B. The False Promise of Protecting Domestic Workers

The screening system for admitting temporary workers creates a conflict of interest that threatens the protection of domestic workers. The conflict arises because the current rules delegate to employers the task of screening prospective temporary labor migrants.³

Such delegation makes good sense from an informational perspective. Firms are generally better equipped than the federal immigration bureaucracy to sift through the millions of applicants living outside the United States. And firms obviously know more about themselves, making them better able to make good matches (Cox & Posner 2013). In an allocation problem rife with externalities, however, it is less clear that putting employers in charge makes sense from society's perspective. Firms will, in principle, focus only on their private benefits and costs, and their incentives need not be well-aligned with the goal of protecting domestic workers from foreign labor market competition. The very existence of numerical limits on labor migration is evidence of this mis-alignment.

To ameliorate this concern, the immigration code has long attempted to build protection for domestic workers into the criteria for approving visa petitions submitted by employers. Under the traditional version of these "labor certification" rules (INA § 212(a)(5)(A)(i)), the employer must show, to the satisfaction of the Department of Labor, that:

- (1) there are not sufficient [domestic] workers who are able, willing, qualified . . . and available at the time of application for a visa and admission to the United States and at the place where the alien is to perform such skilled or unskilled labor, and
- (2) the employment of such alien will not adversely affect the wages and working conditions of workers in the United States similarly employed.

These requirements are quite clearly designed to protect domestic workers. But despite the fact that the basic idea of labor certification has been around since 1952,⁴ the labor market principles underpinning these twin requirements remain obscure. After all, it is difficult to imagine a sustained shortage of workers when the wage can be set freely. Moreover, in the simplest of labor market models, adding additional workers from overseas would always "adversely affect the wages" of existing workers—simply by increasing the labor supply. One might wonder, therefore, how these requirements would ever be satisfied.

In principle one might rationalize the labor certification requirements with additional theorizing. Perhaps, for example, the first requirement is intended to limit use of temporary employment visas to situations where the labor market experiences an exogenous shock. In

³ For an overview of this core feature of immigration law, see Cox & Posner (2012). It is true, of course, that there is also a set of qualitative pre-requisites that migrants must satisfy. To qualify for a H-1B visa, for example, immigration law (8 U.S.C. § 1184(i)) requires that the foreign worker must: (1) come to work in a "specialty occupation," defined as one that requires "technical and practical application of a body of highly specialized knowledge;" and (2) have a bachelor's degree in that specialty. But these screening criteria are generally much less restrictive than those imposed by employers themselves.

⁴ The 1952 Immigration and Nationality Act adopted labor certification in place of the old contract labor laws that had, since 1875, limited at least some forms of labor migration.

such a situation there could be a temporary shortage of workers. But while this might rescue the notion of labor certification as a matter of theory, there is no evidence that it actually works this way in practice. Instead, the process mostly serves to bring workers to the same occupations year after year.

Alternatively, one could imagine that the goal of labor certification is to screen potential hiring decisions, permitting firms to hire foreign workers who are complements to domestic workers but prohibiting them from hiring workers who are substitutes. Both requirements—that there be no domestic worker available, and that hiring the foreign worker not “adversely affect” the wages of domestic workers—could be seen as gestures towards the idea that the foreign worker will be a complement to domestic workers. Again, however, there is little evidence that labor certification actually works in this fashion. The text of the labor certification rules themselves—which require that the foreign workers be compared to domestic workers who are “similarly employed”—belies the notion that the legal rules are designed to weed out immigrants who are substitutes to local workers.

In fact, even the limited oversight provided by the traditional labor certification system has been relaxed for visas like the H-1B. First, for H-1B visas, the government’s review of the employer’s labor certification application is replaced by a rule that permits employers simply to attest to the existence of the labor market facts that justify hiring an H-1B worker. The attestation rule flips the default from one where the firm cannot hire a worker until the Department of Labor approves to one where the firm can hire unless the Department disapproves. Moreover, in the vast majority of cases there is no meaningful review of the attestations by the Department. Thus, firms act in the shadow of *potential* monitoring by the Labor Department, rather than in the face of actual monitoring.

Second, the H-1B rules relax the substantive labor market constraints that are part of the traditional labor certification process. The H-1B visa rules (INA § 212(n)(1)) require only that the employer attest that (1) it will pay a wage that is at least either “the actual wage paid by the employer to all other individuals with similar experience and qualifications,” or the prevailing wage set by the Department of Labor, whichever is greater; and (2) that it will provide working conditions for the H-1B workers that “will not adversely affect the working conditions of workers similarly employed.” The employer need not attest that there are no domestic workers available. In fact, unless H-1B workers make up a significant share of its workforce, it is not even required to attest that it has taken good faith steps to recruit domestic workers before seeking an H-1B visa.⁵

Given all this, it should not come as a surprise that the labor certification process is widely seen as a sham. The basic problem is that the legal rules do nothing to generate incentives for employers to prefer domestic workers. Instead, the labor certification rules

⁵ This additional attestation is required of so-called “H-1B dependent” employers, defined as employers with more than 50 employees for whom H-1B workers constitute more than 15% of its workforce. H-1B dependent employers face additional requirements as well. They must pay a larger fee for each petition, for example, and they must attest that they have not displaced a United States worker “within the period beginning 90 days before and ending 90 days after the date of filing” the visa application (INA § 212(n)(1)(E)(i)). Under the code (INA § 212(n)(3)), some of the extra scrutiny otherwise imposed on H-1B dependent employers is waived if the H-1B employee is paid an annual wage of at least \$60,000 (\$100,000 under legislation recently introduced in Congress (H.R. 115-170)).

attempt to solve the problem of employers' mis-aligned incentives by having government bureaucrats police the firms' hiring decisions. But ex post monitoring has proven ineffective. The labor market is simply too complex, and the Department of Labor ill-equipped, to do much more than prevent flagrant fraud within the system—and even on that score it often fails.

C. The Exploitation of Immigrant Workers

The third concern about temporary worker programs is that they lead to the exploitation of the immigrant workers, a longstanding concern motivated by painful historical experience. America's largest guest worker program was the Bracero program which, launched during WW II, eventually involved some 4.6 million workers, with an average of 200,000 braceros coming each year between 1948 and 1964 (Ngai 2004). The program left behind a sad legacy of worker mistreatment. Today, the possible exploitation of temporary workers continues to shape debates over labor migration policy. And while programs involving low-skilled workers, like the modern H-2A program for agricultural workers, have raised the greatest concerns, workers subject to other programs are vulnerable too. Just recently, for example, several large tech companies have been accused of paying substandard wages to computer engineers admitted through the H-1B program (Preston 2016).

There are two very different reasons to worry about the “exploitation” of temporary workers. The first is that these workers might enter into contracts that we may think, for any number of reasons, should be forbidden. The right to work temporarily in the United States is extremely valuable for many migrants, in part because the wage premium to work in the U.S. can be enormous.⁶ Given how valuable this right is, a worker may agree to terms that are “exploitative” even if voluntary. (When and why voluntary transactions might be considered exploitative is a complex subject we bracket here.) A second concern is that temporary labor visas empower employers to depart from the agreed-upon contract once the immigrant worker enters the country. This is best understood as a commitment problem—the challenge of holding the employer to the terms of the employment contract after the migrant worker enters on the visa.

The legal rules regulating temporary labor admission contribute to this commitment problem. With the worker's visa tied to her sponsoring employer, the foreign worker is deprived of the usual remedy employees have when their employers try to shade on the performance of their contractual obligations—the option of simply going to work for another employer. And because the temporary worker lacks job mobility, the firm who initially sponsors her needs not later compete to keep her. In other words, the labor market becomes monopsonistic for temporary workers once they enter the country.

* * *

⁶ Consider, for example, the estimates by Clemens et al. (2008) of a migrant's wage in the United States versus the migrant's wage in the country of origin. Controlling for workers' qualifications and comparing wages in terms of purchasing power in the country where they are earned, the median ratio over 43 countries is estimated at 2.5, with much higher peaks for some countries—for Nigeria the ratio is more than 8.

These three important problems have prompted calls for reform. To better protect domestic workers, many have proposed either reducing the quotas for particular visa categories or increasing government scrutiny of employers' visa petitions. To improve the efficiency of visa allocation, Simon (1986, 1999), Becker (2010) and others have proposed either taxing immigrants or auctioning visas to them. These proposals harken back to the nineteenth century origins of American immigration policy, when states frequently tried to use prices (explicit or implicit) rather than non-market mechanisms like quotas and queues to regulate immigration policy. And as conceptualized by their authors, these proposals would be just about as radical as a return to those earlier policies. Adopting them would require abandoning nearly all of the core elements of our existing immigration rules. Employers would no longer play a central role in the immigrant screening system. Visas would be easiest to come by for the wealthy. And promising, productive immigrants of modest means would be unable to immigrate, except in instances where they were able to borrow significantly against their future earnings. In what follows, we propose a solution that would fit naturally in our current employer-driven system and would not discriminate among potential immigrants on the basis of their resources. Moreover, we posit that it has the potential to solve simultaneously all three of the problems we identified above.

II. A POTENTIAL SOLUTION

Our proposed solution hinges on three main points. First, granting visas to firms confers on them a valuable asset. It is appropriate to charge them for it. We thus propose to auction labor visas directly to employers. Existing experiences with allocating valuable public rights teach that the auction will be more effective if complemented by the flexibility of an on-going secondary market, with titles to the visas made fully tradable among firms. The first component of our proposal is therefore an auction-cum-market allocation mechanism for new visas.

Second, making visas tradable requires separating the visa from the identification of a specific employee. As is the case now, each visa bundles (a) the firm's right to employ a foreign worker for a specified period of time, with (b) the worker's right to reside and work in the country during that time. Unlike the existing system, however, the rights remain abstract until the employment contract is written: neither the firm nor the worker are identified, and the rights can be applied to any firm who owns the title to a visa and any worker with the required qualifications.

Third, at the moment an employment contract is signed, the two rights originally bundled together in the visa become separate and personalized. If firm F signs a contract with worker W, F maintains the ownership of the right to employ W for the duration of the visa, but W acquires the right of residence and employment for the specified time. Because the rights are now personalized and unbundled, they differ from the original abstract rights and are no longer tradable.

Note that the worker's right of residence and employment implies the right to change employer. If worker W decides to take a job with firm F', firm F' acquires from F the right to employ W for the remainder of the visa duration. It is then appropriate to charge F' for the pro-rate price of a visa on the secondary market, and to credit F with the corresponding amount.

In summary, our proposal exploits market forces to deliver a better allocation of visas, together with a fee exacted on firms, and benefits distributed between the government and domestic and foreign workers.

Our point of departure—auctioning labor visas directly to employers—is closely tied to the functioning of the current system, and respects the informational advantage over the needs and opportunities of the labor market enjoyed by firms. It was first suggested by one of us in work co-authored with Eric Posner (Cox & Posner 2012) and, in its combination with a secondary market, has since been recommended independently by Peri (2012).

The intellectual foundation of an auction-cum-market system such as Peri’s and ours is Coase (1960)’s fundamental insight: the original ownership of the right is distinct from the use of the right, and efficiency can be enhanced by the possibility of trade. The markets for environmental permits are the most successful practical realization of such insight, and their rich international experience is a trove of lessons on desirable and undesirable design features.⁷

But labor visas are importantly different from environmental permits. Our proposal distinguishes itself from previous ideas by recognizing that while visas can initially be conceptualized as abstract property rights, the specificity of an employment contract alters their nature fundamentally. Post-contract, the visa changes from an abstract right to a concrete, individual relationship, and as such should be non-tradable. We argue that the appropriate response to such a change is the separation of the right-of-residence and the right-of-employment. In what follows, we describe more precisely the various components of our proposal.⁸

A. The Visas

For simplicity, we describe our proposal in terms of a single category of visas. Different categories correspond to different types of employment and should be treated separately: sold in separate auctions and traded in separate secondary markets. For each category, the yearly cap is fixed as a matter of law. The quota could be established by statute, as it is today for all capped categories of labor admission. Or the authority to establish the cap could be delegated to the President (or an administrative agency), as is the case today for the annual refugee admissions quota (Rodriguez 2010).

As happens now, each visa would bundle the employer’s right to employ the foreign worker for a fixed period, and the worker’s right to reside in the United States for the same fixed period. The period would likely vary for different visa categories, as it does under current law: H-1B visas, for example, authorize up to a three-year term of employment and

⁷ A useful short summary of the experience of different environmental markets is Ausubel et al. (2014). The implementation of Coase’s insight via appropriate markets for property rights has been proposed to address many different types of externalities. One of us (Casella 1999), for example, has suggested it as a better tool to achieve fiscal discipline among the countries of the European Monetary Union.

⁸ In an online appendix, we provide more detail about various aspects of the implementation—for example the specific characteristics of the visas; the timing, size, frequency, and format of the auctions; the organization of the secondary market. [Electronic address] Our suggestions are preliminary and open to improvement and fine-tuning, but are meant to make clear that the proposal is both concrete and feasible.

admission, while H-2A agricultural visas authorize a maximum term of admission of one year.⁹

Because pre-contract visas are abstract rights, they can be converted into specific employment contracts at different times. The design of desirable market mechanisms is helped if visas do not expire: the year for which they are issued is the earliest date at which they can be used, but firms can store the visas for future foreign employment. If visas lack an expiration date and all visas can be used as soon as they are issued, then all pre-contract visas in circulation are perfect substitutes: all grant the same right to employ a foreign worker for the specified duration of employment, and all command the same market price. Firms' ability to choose freely when to convert the visas they own into employment contracts introduces aggregate flexibility; foreign workers will be brought into the country when the economy is best able to absorb them. Note that flexibility is important at the level of the individual firm too. If a firm faces uncertainty, it can postpone its plans to hire foreign workers and store its visas, without having to transact on the secondary market. Finally, but equally important, if all existing pre-contract visas are perfect substitutes, the market will be more liquid and thick, and the values of the visas more stable and less responsive to events of the specific year in which they were put in circulation.

The absence of an expiration date could in principle lead firms holding visas to lobby for lower quotas in the future, increasing the value of the visas they own and creating a bias towards an increasingly restrictive immigration policy. In practice we doubt this should be a concern. Firms without visas would have the opposite incentive, and lobby for increased quotas in the future in order to depress the value of the visas held by other firms. And even firms holding visas may prefer an increased supply: firms holding, say, valuable oil and gas rights on government lands do not always lobby for greater future restrictions on access to those rights. Whenever property rights constitute both a valuable financial asset and a factor in production for a firm, there is no a priori reason to think that the firm will prefer either more or less scarcity of the property. Often, what we observe is that familiarity and success in navigating regulatory constraints lead to a vested interest in maintaining the set of existing rules relating to those rights.

B. The Auction and the Market

Whenever the government needs to distribute scarce, valuable property rights—whether spectrum licenses or pollution permits, government bonds or procurement contracts—it faces an allocation problem. Temporary visas are no different. As in other contexts, there is first the question of who should initially receive the property rights; and second the question of whether the property rights can later be traded on a secondary market.

⁹ Having a single time length for all visas in the same category makes them perfect substitutes and simplifies and thickens the market. Note that employers who wish to hire a worker for a shorter period will ultimately be compensated for the unused portion of their visa, as long as the worker utilizes the balance of the visa working for another employer in the United States. An alternative, more flexible design would slice the length of the visa into smaller time increments, making it possible to trade, for example, monthly or quarterly permits. Although theoretically possible, the need to obtain usable permits by aggregating bundles of smaller units would complicate the design and the intuitive understanding of the scheme we propose.

These twin questions are linked. If the secondary market is thick and non-distorted, the initial allocation of the rights has little effect on the final allocation, and thus on the efficiency of the system. (It does of course affect the distribution of the gains and losses among the market participants). Logically, it would be conceivable to make pre-contract visas tradable, and yet continue to allocate them to firms as they currently are. In the earliest environmental markets, for example, licenses were freely distributed to regulated firms as a way of reducing industry opposition to the new regulatory regime.¹⁰ This motivation seems weaker in the case of temporary work visas, where the political picture is complicated by the presence of the foreign outsourcing firms. The wind-fall they receive from the current allocation of visas does not seem particularly desirable, either economically or politically.

There are other reasons to recommend an auction to grandfathering. First of all, if the secondary market is distorted, the initial allocation of the visas will matter. The auction will tend to allocate permits to those firms who value them most, and thus enhance efficiency. Second, of course, the auction will raise government revenue, a desirable result both on fiscal grounds and on principle, given the attribution to firms of a valuable social right.¹¹

Who is allowed to take part in the auction? And should supply consist exclusively of new visas sold by the government, or should private resale be allowed? In general, there are good theoretical reasons for making participation in the auction as wide as possible and for allowing private supply. Both features make collusion more difficult to coordinate and enforce, and increase competition and efficiency (Holt, 2007; Milgrom, 2002). On practical grounds, on the other hand, it seems desirable to limit participation to firms, as potential employers, and to restrict supply to the government: the proposal to allocate visas through market mechanisms will raise strong concerns, the stronger the more open to diverse interests the market is perceived to be. There is also an additional reason for limiting participation to firms, specific to our proposal. When a worker employed under a temporary visa chooses to change employment, we suggest using the (pro-rata) visa price in the secondary market as compensation for the incumbent firm and as a fee for the new employer. For the adjustment to be correct, it is important that the market price of visas reflects the *firms'* value for the visa, and no additional considerations.¹² Because all pre-

¹⁰ The classic example is the SO₂ emissions trading program, the first national environmental market. The program was established as Title 4 of the 1990 Clean Air Act and began being operative in 1994. Although considered a major policy success, political interference caused large instability in prices starting in 2005, and the program effectively ended by 2010. For an overview of the program see Ellerman et al. (2002); for an ex-post evaluation see Schmalensee and Stavins (2013).

¹¹ Environmental programs too now increasingly introduce new licenses via auctions, at times after a short initial phase of grandfathering (Ausubel et al., 2014b). For a discussion of the superiority of auctions to grandfathering, see Cramton and Kerr (2002). For experimental results, see Goeree et al (2010) and Olson, Hoffman and Houser (2010). There is general agreement that the optimal design combines an initial auction and a secondary market.

¹² The difficulty is that, when traded, visas include not only the right to employ a foreign worker, but also the worker's right of residence. Such a right passes to the worker automatically and freely upon the signature of the contract (when the visa stops being tradable). It has no positive value to firms and it will not be priced into the market value of the visa as long as visas are purchased by firms alone. If other intermediaries or workers themselves were actively trading, the market value would reflect the separate value of this right, and thus the market price would not be the appropriate compensating reference when workers change employment.

contract visas are perfect substitutes, the market price and the auction price are strictly connected. Our preliminary suggestion then is that both the auction and the secondary market be restricted to potential employers.

Bidders in the auction will typically bid for multiple units. Multi-unit auctions are the allocation mechanism of choice for many public assets, and have been widely studied, tested, and adopted.¹³ Different formats are possible but a large body of theory, experimental studies, and pragmatic experiences has reached two main conclusions. On one hand, simplicity of the rules is essential. It makes implementation easier, encourages participation, and simplifies the bidders' strategic problem. On the other hand, with the exception of very special environments, the simpler multi-unit auction designs do not guarantee that the objects will go to those bidders with highest values. We discuss the source of the problem in the online appendix. Here we just remark that practical multi-unit auctions tend to induce large bidders to underbid, relative to their true values. As a result, the auction's misallocation favors firms with fewer demands for visas, who are more likely to win the auction and yet pay less for the objects than their true market value, and can then resell on the secondary market. Note that favoring smaller firms would counter the perception that the allotment of visas has so far been dominated by larger players.

The possible inefficiency of the auction is mitigated by trading in the secondary market, which again practical considerations suggest limiting to potential employers. The market value of a visa is the value of a foreign worker's employment relative to the available alternative—domestic employment, if domestic labor is available for that job, not filling the position, if not. If the market is fully competitive, the visa price that equates demand and supply reflects the “shadow value” of the immigration constraint—that is, the value of admitting one extra foreign worker, employed in her most productive position. In such a case, the allocation of visas that results from market trading is efficient: a firm is driven to buy visas as long as the value of a foreign worker in its employment is higher than the price, i.e. higher than the best use of the worker in the most productive unfilled position.¹⁴ Contrary to traditional regulations—individual quotas or Pigouvian taxes—there is no need to know all firms' technologies and alternatives, either for the regulator or for the individual traders. The market price summarizes all relevant information.¹⁵

It is important to note that the visa's market price will depend on the expected salary a foreign worker would earn. This implies then that the visa price and the migrant worker's

¹³ For excellent surveys of the theoretical properties of different designs, accompanied by comments on practical implementation, see Ausubel et al. (2014a, 2014b), Milgrom (2004), Klemperer (2004), and Krishna (2002). For specific focus on the inefficiency of multi-unit auctions, see Ausubel and Cramton (2002) and Ausubel (2004).

¹⁴ As stated, the description holds in a static competitive market. In fact, if visas are storable over time they become stores of value, and hence financial assets. Their price will then reflect both their most advantageous immediate use and the expectations of future price increases. In environmental markets, for example, spot markets are now typically accompanied by more complex contracts, including future contracts and options.

¹⁵ When traders are few, or when some of them either own or demand a large share of the visas, the market price will in general be affected by the actions of the large traders. Relative to the competitive model of the theory, in practice distortions are likely, but experience has nevertheless found these markets to be very effective, as long as the supply of permits is not too high and political uncertainty not too disruptive. (Milgrom, 2004; Ellerman et al., 2002).

salary move in opposite directions: the higher the salary, relative to domestic workers, the lower the value of the visa, and the lower its price, and vice versa. Allocating visas through market mechanisms introduces a virtuous circle that reduces firms' incentives to fight higher wages for immigrant workers. Firms' total costs from foreign hires consist of both the visa price and the wage, and with the two components moving in opposite directions, the effect on total costs of higher wages is ambiguous, depending on both foreign workers' labor supply, and the domestic labor market. We will discuss in the next section a simple formal model designed to make these interactions more transparent. For now, notice that given the auction-cum-market allocation scheme, current measures designed to put a floor on foreign wages—which we imagine would be retained under our proposal—should encounter lower resistance from employers.¹⁶

The same interaction between foreign workers' wages and the visa's market price also means that distortions in one market can be compensated in the other. If firms maintained monopsonistic power over hired foreign workers, they could pass on the price of the visa to the worker, in the form of a reduced salary. But firms' ability to do so would increase the value of the visa, and hence its market price, until competition on the visa market would cancel the excess profits. Similarly, if less likely, competition in the market for foreign workers would translate into higher wages if firms were able to appropriate the visas at costs lower than their value. As the formal model will make clear, the two markets are linked and jointly determine the allocation of foreign employment.

A recent proposal by President Trump's transition team advocated awarding visas to employers who commit to pay the highest wages (Rosenberg 2017). This would be an alternative allocation method to our market-based tools, designed with the explicit goal of increasing foreign workers' wages. In general, it is preferable to choose instruments that are directly tied to the goal intended—support wages by increasing competition in the market for foreign workers, as our proposed scheme would do, rather than via the visas' allocation scheme—because unforeseen consequences are more difficult to avoid when the channels are indirect and complex. But even on its own terms the proposal is unlikely to work: it would shift visas towards sectors and employments with highest wages, without necessarily reducing the competitive pressure on domestic workers within those sectors. At the end, because wages differ both across sectors and across geographical regions, the proposal would have to create a profusion of visas, differentiated by sectors and areas. And even if enforceable and successful, it would allocate all of the surplus to workers—to migrant workers, directly, and to domestic workers via the reduced competition—rather than permitting the government to capture some of it.

C. Changing Jobs

Where our proposed scheme distinguishes itself fundamentally—both from existing cap-and-trade programs and from proposed market-based reforms to immigration—is in the

¹⁶ Current law (8 U.S.C. § 1182(n)(1)(A)(i)) already requires employers to pay either the same wage paid to domestic employees or a government-listed “prevailing wages” for the occupation at issue. Moreover, existing H-1B visa rules ((8 U.S.C. § 1182(n)(3)(B)) create incentives for employers to pay H-1B workers at least \$60,000 per year by exempting them from various regulatory obligations when they pay at least this minimum wage.

treatment of the visas post-employment contract. The difference seems inevitable, given the nature of the rights being traded: a person's right to live and work in the country. In our scheme, upon signing the contract, the employee becomes a legal temporary resident, and remains so for the duration of the visa, even if the employment relationship is severed. Such right of residence is not tradable.¹⁷

The logical distinction between tradable pre-contract visas, and non-tradable post-contract rights seems to us inescapable. If residence rights remained bundled and the visa remained tradeable after an overseas worker entered the country and began working for a firm, then any trade by the firm would constitute a trade to another firm of the right to employ that particular worker.

There are at least two serious problems with such a system. First, from the perspective of the employees, it would mean an unacceptable loss of control over the choice of employer. At a minimum, then, any trade would need to be approved by both worker and firm—though the requirement of joint assent would still, as the current system does, leave the worker without defenses against possible abuses. Second, from the firms' perspectives, trades of permits attached to specific workers suffer from serious problems of asymmetric information. Because the original employer has more information than the acquiring firm about the worker's productivity, firms will be reluctant to purchase permits that have become embodied in specific workers: the sale itself is a signal of low productivity.¹⁸

Our recommendation then is that, post-contract, the right of residence and employment be fully vested in the specific worker and not alienable. As owner of such a right, the worker will be free to change employment: she has the right to solicit work from other firms, and other firms are authorized to make her employment offers. The right of employment is fully portable across employers.

Notice that if portability is unconditional, the new employer need not be in possession of the necessary visa, and thus would be able to hire a foreign worker already in the country at a lower cost than if the worker were brought in from abroad. In addition, the original employer would lose part of its investment in the original visa. The distortion could be costly: it would discourage the search and hiring of new foreign workers and incentivize poaching of other firms' workers. To prevent this from happening, the new employer will be required to pay to the government a fee equal to the cost of a visa on the secondary market at the time of hiring, corrected for the proportion of time the worker is still entitled to spend in the country. At the same time, the old employer will be compensated in identical amount, in the form of a credit from the government. We prefer government fees and credits to

¹⁷ The secure right of residence held by the noncitizen should, as a side-benefit, encourage country-specific investments that may be foregone in contexts where the noncitizen's right is contingent on her continued employment by the sponsoring firm (Cox and Posner, 2007).

¹⁸ As in Akerlof's (1970) "market for lemons." There are interesting parallels here to sport markets. But in sports, the level of a player's performance—how well the player plays—is not contractible; the lack of possible enforcement *de facto* gives the player strong bargaining power, and the ability to force a trade the player desires. In addition, the asymmetric information problem is less relevant to sports, because the player's performance in competition is public.

direct transfers between the two firms to prevent either firm from attempting to extract excess compensation from the transaction.¹⁹

Of course the visa's portability will help the worker only to the extent that there is demand for the worker's services. If the skills firms need are highly specialized, workers brought into the country by other firms need not be the best hires, and searching directly on the foreign market may be preferable. In such a case, the demand for a worker already in the country, originally hired by a different firm, may be low, and the worker's bargaining power relative to her original employer will be weak. Although possible, the practical importance of such a concern is likely to be small. Lower-skill jobs, the ones most in need of safeguards, are hardly specialized. Moreover, workers already in the country, somewhat familiar with American ways and beyond the first obstacles of integration, probably make more rather than less desirable employees.

Granting a foreign worker automatic right of residence, independently of current employment, will raise concerns of abuse. What if the migrant loses or leaves her job? In principle, making the right of residence fully unconditional on employment is appealing, because it simplifies the structure of the right and removes the need for the government to monitor continuing employment by the migrant. Broader fiscal consequences are limited by the fact that temporary migrants are ineligible for most forms of public assistance.²⁰ That said, imposing a requirement of continued employment (which exists in current law already) seems desirable in practice and an appropriate incentive compatibility provision in theory: the visa might provide that an immigrant who is unemployed for a period of 90 days or more, for example, forfeits her right to reside and must return to her home country.

Some concern about strategic behavior is not unjustified: immigrants who have no interest in working in the US but would like to live here could be induced to pay employers to purchase a visa on their behalf.²¹ Indeed the possibility of such side payments already exists in our current temporary labor system. It is inherent in any scheme that delegates to employers the power to select visa recipients. Perhaps for that reason, existing immigration rules (INA § 212(n)(2)(C)(vi)) make it unlawful for an employer to accept any reimbursement or compensation from the visa recipient sponsored by the employer. We propose to continue such prohibition. But note that in a system where visas are auctioned rather than given away to employers for free, employers' incentive to demand side payments should be reduced.²²

¹⁹ Alternatively, the acquisition and sale of the necessary firm permits could take place via the secondary market. However, that would require making the permits fully divisible (in terms of time units), to account for the pro-rata permits being transacted.

²⁰ At the federal level, immigrant access to public assistance was dramatically limited by the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (P.L. 104-193).

²¹ This would make the system more like a direct auction of visas to migrants, reducing its sensitivity to domestic labor market conditions.

²² To see why allocating the visas via market mechanisms helps to alleviate the problem of side payments, consider the two effects of side payments in such a system. First, side-payments from workers effectively subsidize the firms' purchase of visas, increasing the firms' demand, and thus the visas' price. To the extent that this happens, the side-payment is effectively transferred to the government, decreasing the employer's incentive

III. EFFECTS ON DOMESTIC FIRMS, DOMESTIC WORKERS AND MIGRANT WORKERS

What would be the consequences—for firms, domestic workers, and migrant workers—of adopting the regime we just sketched out? Two components drive its effects: the market-based allocation of the pre-contract visas (via the auction and the secondary market), and the transfer of the right to reside and work to the migrant post-contract.

The first component creates incentives that favor shifting visas, and foreign employment, towards firms where such employment is most productive. The result is an increase in efficiency—a higher total value of production, for given amount of resources. In addition, firms will have to pay for a right that now accrues to them at very low cost, and thus there will be redistribution from the firms to the government. If the visa price fully reflects its value, all firms' surplus is transferred to the government; if there are distortions and the price does not reflect the visa's full value, then the surplus is shared between the government and the firms able to secure visas below valuation.

The extent and distribution of the surplus will also depend on the second component of our proposal, the post-contract transfer of the property right to the migrant. Transferring the right induces competition for foreign workers and increases their salaries (or more generally improves their working conditions). Such competition is a second source of efficiency gains for the economy as a whole, both because the higher wage causes further pressure towards the most productive use of foreign workers, and because it reduces the dead-weight losses associated with monopsony (the strategic underutilization of resources). At the same time, competition reduces profits from foreign hiring, and thus the market price of the pre-contract visas. In principle, the decline in the visa price could be either bigger or smaller than the increase in the wages of temporary workers. Thus, this second redistributive channel has ambiguous effects on firms' profits—though it clearly favors foreign workers and reduces government revenue from the visa auction.²³

The impact on domestic employment and wages depends crucially on technology—in particular, on whether domestic and foreign labor are complements or substitutes. Competition, both for visas and for foreign workers already in the country, will redistribute foreign employment towards firms and occupations where foreign workers are more productive. If such occupations are ones where foreign workers are primarily complements to domestic workers, then the demand for domestic labor will increase, and with it the pressure on domestic wages. If the shift is instead disproportionately towards sectors where foreign and domestic workers are substitutes, however, then both domestic employment and wages will be negatively affected.

to seek such a payment. Second, in our scheme firms employing foreign workers already in the country are required to pay the visas' market price (pro-rated to the time remaining on the worker's visa). In the presence of side-payments to the firms that first bring the migrants to the United States, the visa market price would be distorted (would be too high) from the point of view of the second employer. This would discourage re-employment and would depress the foreign workers' wages, further reducing the incentive for side-payments from the point of view of the foreign worker.

²³ Of course, in some circumstances the second channel will produce no effects. This could be the case if, for example, competitive forces did not raise foreign wages above already existing, legally imposed “prevailing wage” floors.

As this short description makes clear, the channels are multiple and interdependent, and the final effects depend on features of the economy and of technology that are highly disputed empirically. The model is not meant to answer empirical questions; it is a transparent expository device that clarifies the logical links between the different parts of our proposal.

A. A market for work visas: a simple model

Consider a simple economy where labor is the only factor of production. Labor consists of domestic labor, l_d , and foreign labor, l_f . The economy is composed of two sectors, S and C . In sector S , domestic and foreign labor are perfect substitutes—what matters is total employment, not its composition. Output is produced via a concave production function $f_S(l_d + l_f)$. Because specific examples make the workings of the model more transparent, suppose:

$$f_S(l_d + l_f) = (l_d + l_f)^{1/2} \quad (1)$$

In the C sector, on the other hand, domestic and foreign labor are complements. An increase in one type of labor makes the other type more productive:

$$f_C(l_d, l_f) = (l_d l_f)^{1/2} \quad (2)$$

The number of firms in each sector is fixed, and we denote it by n_S , in the S sector, and n_C in the C sector. The total number of firms in the economy is $n = n_S + n_C$. Both sectors' outputs are sold on competitive markets and both are sold at unit price.

The domestic labor market is also fully competitive, and each firm faces a given wage for domestic labor equal to w . We capture impacts on domestic workers through changes in employment levels.²⁴

The foreign labor market is instead monopsonistic: in line with recent literature,²⁵ we capture the subordination of a foreign worker's legal status to his employer as granting the employer monopsonistic power. Each firm i faces an upward-sloping labor supply for foreign workers $h_{if}(w_f)$. We suppose:

$$l_{if} = h_{if}(w_f) = h_f(w_f) = b w_f \quad (3)$$

where w_f is the wage paid by the firm to its foreign workers, and b is a fixed positive parameter. Each firm hiring foreign workers faces the same labor supply (in both sectors), and we drop the superscript i .

²⁴ In other words, we assume that the impact of the visa program on economy-wide domestic wages is small. Rather, it works through domestic employment in the affected sectors.

²⁵ See for example the discussion in S. Naidu, Y. Nyarko, and S-Y. Wang (2016).

B. *Unconstrained profit maximization*

Suppose first there is no constraint on firms' ability to hire foreign workers. What is the profit maximizing allocation of labor and production? In each sector $J = S, C$, each firm solves the problem:

$$\max_{\{l_d, l_f\}} f_J(l_d, l_f) - wl_d - w_f l_f$$

where, by equation (3), $w_f = l_f/b$.

In sector S , a firm's first order conditions are given by:

$$(1/2)(l_{sd} + l_{sf})^{-(1/2)} = w \quad (4)$$

$$(1/2)(l_{sd} + l_{sf})^{-(1/2)} = 2l_{sf}/b \quad (5)$$

Substituting the first equation into the second, we obtain:

$$l_{sf} = bw/2 \quad (6)$$

and thus:

$$w_{sf} = w/2 \quad (7)$$

$$l_{sd} = 1/(4w^2) - l_{sf} = 1/(4w^2) - bw/2 \quad (8)$$

where the non-negativity constraint on l_{sd} requires $bw/2 \leq 1/(4w^2)$.

In the S sector, foreign labor is hired up to the point where its wage equals half the domestic wage. Profit maximization pins down total employment in the sector, and domestic employment in the sector is residual. When constraints on free hiring of foreign labor are introduced, they will be matched by exactly matching readjustments in domestic labor.²⁶

Profits then equal:

$$\Pi_S = 1/(4w) + bw^2/4 \quad (9)$$

In sector C , the first order conditions are:

²⁶ Note that when l_d and l_f are perfect substitutes, they have equal marginal product (up to a multiplicative constant which here we set to 1) and thus, for arbitrary production function, if the domestic labor market is competitive, foreign labor's marginal product must equal w . With foreign labor supply $h(w_f)$, in the absence of hiring constraints, the following condition must hold: $w = w_f(1 + 1/\eta(l_f))$ where $\eta(l_f)$ is the elasticity of the foreign labor supply at the profit maximizing level of l_f : $\eta(l_f) = (dh(w_f)/dw_f)(w_f/h(w_f))$. In our specification, with linear $h(w_f)$, $\eta = 1$ everywhere. The firms' monopsonistic power over foreign labor allows them to pay foreign workers at half their marginal product.

$$\sqrt{\frac{l_{cf}}{l_{cd}}} = 2w \quad (10)$$

$$\sqrt{\frac{l_{cd}}{l_{cf}}} = \frac{4l_{cf}}{b} \quad (11)$$

which lead to:

$$l_{cf} = b/(8w) \quad (12)$$

$$w_{cf} = 1/(8w) \quad (13)$$

$$l_{cd} = l_{cf}/(4w^2) = b/(32w^3) \quad (14)$$

Note that in the C sector the demand for domestic labor increases with foreign employment because the two factors of production are complements (as opposed to decreasing with foreign employment, as in the S sector, where the two factors are substitutes).

Profits are given by:

$$\Pi_C = b/(64w^2) \quad (15)$$

In sector S , at the optimal allocation the marginal product of foreign labor must, as we explain in footnote __, equal w . In sector C , the first order conditions together with equation (3) imply that at the optimal the marginal product of foreign labor must equal $1/(4w)$. It follows that in this model the relative productivity of foreign labor in the two sectors is pinned down by the value of w : if $w > 1/2$, l_f is more productive in the S sector—an intuitive result since substituting away from domestic labor must be more advantageous at higher domestic wages. If instead, $w < 1/2$, l_f is more productive in the C sector, where it works to increase the productivity, and hence the hiring, of domestic labor. The different productivity is mirrored in the other endogenous variables: $l_{sf} > l_{cf}$ and $\Pi_S > \Pi_C$ if $w > 1/2$, and $l_{sf} < l_{cf}$ and $\Pi_S < \Pi_C$ if $w < 1/2$.

C. Constraining total foreign employment

Suppose now that a ceiling \bar{L}_f is set on total foreign employment. The question we ask is how different mechanisms for allocating visas under this ceiling affects the labor market and firms' profit. To compare our proposal to existing immigration policy, we model three different scenarios: first, one where each firm receives the same number of visas, which we call the case of "equal quotas;" second, a scenario under which visas are allocated by a market in which firms maintain ownership of the visas; and third, one in which the visa market is supplemented by a rule that shifts the right of residence to the worker when the employment contract is signed. The equal quotas case is meant as an approximation of the expected allocation under today's (effective) lottery system. It ignores both ex post differences in realized allocation under the lottery (random, and a source of further inefficiency), as well as firms' ex ante strategic positioning in the lottery (which may or may

not favor more productive firms—recall the discussion of foreign outsourcing firms—and is wasteful in itself).

1. *Equal quotas (approximation of lottery)*

We suppose that the quotas are everywhere binding: $\bar{L}_F/n \leq \min(bw/2, b/(8w))$, and thus $l_{Jf} = \bar{L}_f/n$, for $J = S, C$. With $l_{Sf}^q = l_{Cf}^q = \bar{L}_f/n$ (where the q superscript stands for “quotas”), foreign wages are equal in the two sectors:

$$w_{Sf}^q = w_{Cf}^q = \bar{L}_f/bn \quad (16)$$

The values of the other variables is derived immediately from our previous equations:

$$l_{Sd}^q = 1/(4w^2) - \bar{L}_f/n \quad (17)$$

$$\Pi_S^q = 1/(4w) + (\bar{L}_f/n)(w - \bar{L}_f/(bn)) \quad (18)$$

$$l_{Cd}^q = 1/(4w^2)(\bar{L}_f/n) \quad (19)$$

$$\Pi_C^q = 1/(4w)(\bar{L}_f/n) - (\bar{L}_f/n)(\bar{L}_f/bn) \quad (20)$$

The effects of the quotas are straightforward: relative to the unconstrained case, the decline in foreign employment in both sectors results in higher domestic employment in the S sector (where domestic and foreign employments are substitutes), and lower domestic employment in the C sector (where they are complements). Foreign wages fall in both sectors, as do profits. The relative decline in foreign employment and in profits across the two sectors depends on the relative severity of the constraint, and thus on foreign labor productivity in the unconstrained optimum. If foreign labor’s productivity is higher in the S sector (equivalently, in our model, if $w > 1/2$), then the decline in foreign employment, foreign wages, and profits is larger in the S sector; if instead foreign labor’s productivity is higher in the C sector (or $w < 1/2$), then the opposite holds.

2. *A market for visas 1. The employer has the final property right*

Consider now implementing the same aggregate ceiling \bar{L}_f via a market for visas. The visas are introduced in the system by an auction, but if the secondary market clears at price p , then p must also be the winning bid at the auction. What matters for the firms is that the unit cost of a foreign worker is now $(w_f + p)$. We want to solve for the market-clearing price p , for the allocation of labor, domestic and foreign, and for the sectoral profits that follow from this scheme.

We assume that only firms can trade on the market, that they can trade only before a specific employment contract is signed, and that the visa market is perfectly competitive: each firm takes the price p as given. There is a single visa market for the economy: firms from both sectors buy their visas on that market.

The market lends flexibility to the allocation of foreign employment across the two sectors. Qualitatively, the impact is clear: relative to equal quotas, foreign employment will be higher in the sector that the quota would restrict more severely, i.e. the sector where foreign labor is more productive in the absence of constraints. If foreign labor is more productive in the S sector, then allocating visas via market mechanisms will allow higher foreign employment and higher foreign wages in that sector, relative to quotas, and lower foreign employment and lower foreign wages in the C sector. Domestic employment will, relative to equal quotas, be lower in both sectors. If instead foreign labor is more productive in the C sector, then all effects are reversed, and domestic employment is higher in both sectors.

Asking firms to buy visas, as opposed to receiving them for free, reduces firms' profit. But here too the effect is mediated by the reallocation of foreign labor that market mechanisms make possible. Relative to quotas, the decline in profits is smaller in the sector where foreign labor is more productive. In fact, if the difference in productivity is large enough, and thus the value of the labor reallocation high enough, in principle profits could be higher than under quotas in the high productivity sector, the cost of the visa more than compensated by the flexibility provided by the market.

Implementing the aggregate ceiling via a visa market (cum auction) is superior to quotas in terms of overall efficiency (total firms profits plus government revenues must be higher than total profits under quotas). Relative to quotas, the scheme also acts as a redistributive mechanism from the firms to the government.

Solving the firms' problem when the unit cost of a foreign worker is $(w_f + p)$, we obtain the firms' optimal employment levels in the two sectors as function of p :

$$l_{Sf}^m = \frac{b}{2}(w - p) \quad (21)$$

$$l_{Cl}^m = \frac{b}{2}\left(\frac{1}{4w} - p\right), \quad (22)$$

where we use the superscript m to indicate the existence of a market for visas. The visa market-clearing price p must be such that the total demand for visas equals the available total supply \bar{L}_f , if $p > 0$, or $p = 0$ if demand falls short of supply. Thus p solves:

$$p[n_s l_{Sf}^m + n_c l_{Cl}^m - \bar{L}_f] = p\left[\frac{n_s b}{2}(w - p) + \frac{n_c b}{2}\left(\frac{1}{4w} - p\right) - \bar{L}_f\right] = 0 \quad (23)$$

or:

$$p = \frac{1}{n}(n_s w + \frac{n_c}{4w} - \frac{2\bar{L}_f}{b}) \quad (24)$$

if $n_s w + n_c/(4w) > 2\bar{L}_f/b$, a condition that is always satisfied under our previous assumption that equal quotas are always binding: $\bar{L}_f/b \leq \min((bw)/2, b/(8w))$.

Substituting the equilibrium visa price into the firms' labor demands, we obtain:

$$l_{Sf}^m = \frac{\bar{L}_f}{n} + [w - \frac{1}{n} (n_s w + \frac{n_c}{4w})] \quad (25)$$

$$l_{Cf}^m = \frac{\bar{L}_f}{n} + [\frac{1}{4w} - \frac{1}{n} (n_s w + \frac{n_c}{4w})] \quad (26)$$

The equations yield the conclusion described above: if the productivity of foreign labor is equal in the two sectors ($w = 1/2$), the visa market does not lead to any reallocation of labor, relative to the case of equal quotas: $l_{Sf}^m = l_{Cf}^m = \bar{L}_f/n$. But the market has the flexibility to respond to differences in productivity: $l_{Sf}^m > \bar{L}_f/n > l_{Cf}^m$ if $w > 1/2$, and $l_{Cf}^m > \bar{L}_f/n > l_{Sf}^m$ if $w < 1/2$.

Foreign workers' wages in the two sectors mirror the sectors' labor demands ($w_{Sf}^m = l_{Sf}^m/b$; $w_{Cf}^m = l_{Cf}^m/b$), and thus will generally differ: because of the firms' monopsonistic power, wages are determined at the individual firm level, and disparity in wages is to be expected.²⁷

Profits are given by:

$$\Pi_S^m = 1/4w + l_{Sf}^m (w - w_{Sf}^m - p) \quad (27)$$

$$\Pi_C^m = l_{Cf}^m/4w + l_{Cf}^m (w_{Cf}^m + p) \quad (28)$$

At $w = 1/2$, $l_{Sf}^m = l_{Cf}^m = l_{Jf}^q = \bar{L}_f/n$, and $w_{Sf}^m = w_{Cf}^m = w_{Jf}^q$, but, under our assumption that the equal quota is binding ($2\bar{L}_f/n < bw/2$), $p = w - 2\bar{L}_f/(bn) > 0$. Hence firms' profits must be lower than under quotas: if $w = 1/2$, there is no difference in foreign labor productivity in the two sectors and no efficiency gain from possible reallocations, while visas have positive costs. If $w \neq 1/2$, however, the decline in profits is mitigated by the labor reallocation, and is always smaller in the sector in which foreign labor is more productive.

The total sum of the firms' profits ($n_s \Pi_{Sf}^m + n_c \Pi_{Cf}^m$) and the government's revenues from the visas ($p\bar{L}_f$) is equal to the firms' total profits under equal quotas when $w = 1/2$, but is otherwise always higher, reflecting the gain in efficiency.

Note one obvious but important point: the visas' market price is determined in equilibrium, reflecting demand and supply forces. Without taking into account how labor productivity in the two sectors affects the demand for visas, and hence their price, we cannot evaluate the final effects on relative employment, wages, and profits. This becomes very relevant when studying the effect of shifting the property right on the visa to the worker, and thus increasing competition for foreign labor.

²⁷ Dispersion in wages for equal workers is indeed considered a tell-tale sign of monopsonistic power.

3. *A market for visas 2. The foreign worker acquires the final property right*

Suppose now that the property right on the visa shifts to the foreign worker upon the signature of the employment contract: the worker can now change jobs, and the monopsonistic power of individual employers is reduced. We model the employers' power as intermediate between the assumption of monopsony we have used so far and full competition, where the foreign worker's wage would equal her marginal product. Again following the literature (Naidu et al. 2016), we assume Cournot competition: all firms, in both sectors, compete for the same foreign workers, facing the same upward-sloping labor supply curve:

$$L_{Sf}^c + L_{Cf}^c = bw_f \quad (29)$$

where $L_{Jf}^c = n_J l_{Jf}^c$ is total foreign labor employed in sector J and the superscript c indicates the presence of competition for foreign workers. Note that w_f is no longer distinguished by sector.

We study a quasi-symmetric equilibrium where each firm within a given sector has the same foreign labor demand, but demands differ across sectors. If non-zero, each firm's foreign labor demand is then given by:

$$l_{Sf}^c = \frac{1}{n_S+1} [b(w - p^c) - n_C l_{Cf}^c] \quad (30)$$

$$l_{Cf}^c = \frac{1}{n_C+1} [b\left(\frac{1}{4w} - p^c\right) - n_S l_{Sf}^c] \quad (31)$$

If the visa price is non-zero—that is, if the constraint on the availability of foreign labor is binding and $n_S l_{Sf}^c + n_C l_{Cf}^c < \bar{L}_f$ —then p is given by:

$$p^c = \frac{1}{n} \left(n_S w + \frac{n_C}{4w} - \frac{(n+1)\bar{L}_f}{b} \right) \quad (32)$$

Equating foreign labor supply with firms' foreign labor demand shows that, as expected, competition increases the wage paid to foreign workers. But the unit cost to firms, per foreign worker employed, must include the price of the visa. Whether $w_f^c + p^c$ is larger or smaller than $w_f^m + p^m$, and by how much, depends on the elasticity of foreign labor supply. With our specification, the equations above yield

$$w_f^c + p^c = w - \frac{\bar{L}_f}{nb} = w_f^m + p^m \quad (33)$$

The unit cost of foreign labor, including the permit price, is *identical* to the unit cost in the absence of competition: the decline in the visa price exactly compensates for the increase in foreign wages. With a linear foreign labor supply (i.e. unitary elasticity of labor supply), the loss of monopsonistic power on the part of the firms translates into redistribution exclusively from the domestic government to the foreign workers, without a direct effect on the firms. In general, whether the decline in the visa price overcompensates or falls short of the increase in foreign wages depends on the elasticity of labor supply. The

more elastic the foreign labor supply, the more the visa price declines relative to the increase in the foreign wage. Thus, if the foreign labor supply is sufficiently elastic, competition for foreign workers already in the country leads to declines in unit costs and, therefore, improves the position of employers (at the cost of the government, which ends up with less auction revenue).

If the productivity of foreign labor differs in the two sectors, competition in the labor market introduces a further channel towards a more efficient distribution of foreign labor: because competition equalizes foreign wages across the two sectors, foreign labor is further reallocated towards the more productive sector. The impact on profits depends not only on unit costs but also on this additional reallocation of labor: relative to maintaining the visa's property right with the firm, foreign employment and profits increase in the sector where foreign labor is more productive (in the S sector if $w > 1/2$, and in the C sector if $w < 1/2$), and fall where it is less productive. Thus there is redistribution among firms from the less productive to the more productive sector. As before, the reallocation of foreign labor also determines the final effect on domestic employment, reducing it if $w > 1/2$ and foreign labor is more productively employed in the S sector, increasing it otherwise.

V. CONCLUSION

In an ideal world, the rules we use to allocate a fixed number of temporary labor visas each year would distribute those visas efficiently (and maximize social surplus), minimize the negative impact on domestic workers, and protect the migrant workers from exploitation. In practice, our current rules are criticized for doing a poor job on all of those dimensions. Our aim has been to highlight the way that temporary labor visas bundle multiple property rights, and to explore whether unbundling those rights and allocating them through market mechanisms could improve the existing system along all three dimensions.

As the model makes clear, this is possible: under our proposal, there could be a virtuous circle where the government gains, foreign workers gain, and domestic employment increases, with costs to firms mitigated by better labor allocation. And, in fact, under nearly all conditions the system we propose would increase both the efficiency with which visas are allocated and the position of temporary workers in the labor market.

Our proposal would increase the bargaining power of temporary foreign workers, and this aspect alone should alleviate competition and act to protect domestic employment. That said, the enhanced efficiency in labor allocation could potentially have negative effects on domestic workers, relative to the existing lottery system. Whether domestic workers ultimately benefit depends crucially on what sectors are more severely constrained in their demand for foreign workers, and what technology these sectors have. If the more severely constrained sectors are those where foreign and domestic workers are complements, then our proposal results at the same time in better allocation of resources and gains for domestic workers; if instead the more severely constrained sectors are those where foreign and domestic workers are substitutes, then the efficiency gains remain, but domestic employment would fall.

The question ultimately is empirical. Still, however it is answered, for a given ceiling on foreign temporary visas, our proposal suggests mechanisms that lead to a more efficient

allocation of labor and better protection. It is difficult to think that allocating the same number of foreign workers inefficiently across firms and sectors is a good long-term answer to the aim of protecting domestic employment.

Table 1. Most Numerically Significant Temporary Labor Admissions Categories

Class of Admission	Annual Visa Quota	Number of Visas Issued FY 2014
Temporary visitors for business (B1/B2)	none	7.6 million
Workers in specialty occupations (H-1B)	85,000	161,369 ²⁸
Low-skilled workers (H-2)		157,376
Agricultural workers (H-2A)	none	89,274
Non-agricultural workers (H-2B)	66,000	68,102
Intracompany transferees (managers and executives) (L1)	none	71,513
Trader or investor authorized by bi-lateral treaty (E)	none	44,155
Internationally recognized athletes or entertainers	none	23,335
Workers of extraordinary ability (O)	none	12,706

Notes: The Immigration and Nationality Act is the source of the visa quotas. Visa issuance numbers are drawn from Department of State, Report of the Visa Office 2014 (Table XVI).²⁹

²⁸ The original statutory quota for H-1B visa was 65,000, but exceptions were progressively added. First, Congress made the first 20,000 H-1B workers with at least a master’s degree from a United States institute of higher education exempt from the H-1B cap, effectively increasing the quota to 85,000. Second, Congress exempted from the cap *any* H-1B worker employed in the United States by an institution of higher education or a related nonprofit, or employed at a nonprofit or government “research organization.” As a result, in 2014 161,369 H-1B visas were issued.

²⁹ For all classes of admission above except the E visa, the visa issuance numbers *do not* include visas issued to the spouse or children accompanying the primary visa recipient. We report the number of visas issued each year, rather than the number of admissions to the country (a more common statistic that is frequently reported) because the government’s admissions data counts each admission into the United States, rather than counting each person admitted. Thus, these data double-count any person who enters the country more than once on the same visa. This can make admissions numbers a relatively uninformative measure for temporary visa categories in which multiple admissions are common.

REFERENCES

- Akerlof, George. 1970. The Market for “Lemons”: Quality, Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* 84:488-500.
- Ausubel, L. and P. Cramton. 2002. Demand Reduction and Inefficiency in Multi-Unit Auctions. Working Paper. University of Maryland.
<http://drum.lib.umd.edu/bitstream/handle/1903/7062/98wp-demand-reduction.pdf?sequence=1>.
- Ausubel, L. 2004. An Efficient Ascending-Bid Auction for Multiple Objects. *American Economic Review* 94:1452-1475.
- Ausubel, L., P. Cramton, C. Aperjiis, and D. Hauser. 2014a. Pilot Auction Facility for Methane and Climate Change Mitigation: Relevant Auction Theory. Power Auctions LLC:
<http://www.pilotauctionfacility.org/sites/paf/files/Review%20of%20Relevant%20Auction%20Theory.pdf>
- Ausubel, L., P. Cramton, C. Aperjiis, and D. Hauser. 2014b. Pilot Auction Facility for Methane and Climate Change Mitigation: Relevant Environmental Auctions, Power Auctions LLC:
<http://www.pilotauctionfacility.org/sites/paf/files/Review%20of%20Environmental%20Auctions.pdf>
- Becker, Gary. 2011. *The Challenge of Immigration—A Radical Solution*. London, UK: Institute of Economic Affairs.
- Burtraw, D., J. Goeree, C. Holt, E. Myers, K. Palmer, and W. Shobe. 2011. Price Discovery in Emissions Permits Auctions. Pp. 11-36 in vol. 14 of *Experiments on Energy, the Environment, and Sustainability (Research in Experimental Economics)*, edited by M. Isaac and D. Norton. Emerald Group Publishing Limited.
<http://www.emeraldinsight.com/doi/abs/10.1108/S0193-2306%282011%290000014004>
- Casella, Alessandra. 1999. Tradable Deficit Permits: Efficient Implementation of the Stability Pact in the European Monetary Union. *Economic Policy* 29:323-361.
- Howard F. Chang. 1997. Liberalized Immigration as Free Trade: Economic Welfare and the Optimal Immigration Policy. *University of Pennsylvania Law Review* 145:1148
- Clemens, Michael, Claudio Montenegro, and Lant Pritchett. 2009. The place premium: wage differences for identical workers across the US border. Working Paper. HKS Faculty Research Working Paper Series RWP09-004, John F. Kennedy School of Government, Harvard University.
- Coase, R. 1960. The Problem of Social Cost. *The Journal of Law and Economics* 3:1-44.
- Cox, Adam B. and Eric A. Posner. 2007. The Second-Order Structure of Immigration Law. *Stanford Law Review* 59:809-856.

- Cox, Adam B., and Eric A. Posner. 2012. Delegation and Immigration Law. *University of Chicago Law Review* 79:1285-1349.
- Cramton, P. 2007. Comments on the RGGI Market Design. <http://www.cramton.umd.edu/papers2005-2009/cramton-rggi-market-design-comments.pdf>
- Cramton, P. and S. Kerr. 2002. Tradeable Carbon Permit Auctions: How and Why to Auction not Grandfather. *Energy Policy* 30:333-345.
- Ellerman, D., P. Joskow, R. Schmalensee, J.-P. Montero, and E. Bailey. 2002. *Markets for Clean Air: The U.S. Acid Rain Program*. Cambridge, U.K.: Cambridge University Press.
- Goeree, J., C. Holt, K. Palmer, W. Shobe, and D. Burtraw. 2010. An Experimental Study of Auctions versus Grandfathering to Assign Pollution Permits. *Journal of the European Economic Association* 8:514-525.
- Holt, C., W. Shobe, D. Burtraw, K. Palmer and J. Goeree. 2007. Auction Design for Selling CO2 Emission Allowances Under the Regional Greenhouse Gas Initiative, Final Report: http://econ.ccps.virginia.edu/RePEc_docs/ceps_docs/rggi_auction_final.pdf.
- Kagel, J. H., and D. Levin. 2001. Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions. *Econometrica* 69:413-454.
- Klemperer, P. 2004. *Auctions: Theory and Practice*. Princeton, NJ: Princeton University Press.
- Krishna, V. 2002. *Auction Theory*. San Diego, CA: Academic Press.
- Milgrom, P. 2004. *Putting Auction Theory to Work*. New York: Cambridge University Press.
- Naidu, S., N Nyarko, and S-Y Wang. 2016. Monopsony Power in Migrant Labor Markets: Evidence from the United Arab Emirates. *Journal of Political Economy* 124: 1735-1792.
- Ngai, Mae. 2004. *Impossible Subjects. Illegal Aliens and the Making of Modern America*. Princeton, NJ: Princeton University Press.
- O'Brien, Sara A. 2016. High-Skilled Visa Applications Hit Record High – Again. *CNN Money*, Apr. 12.
- Olson, M., K. Hoffman, and D Houser. 2010. An experimental investigation of repeated auctions and secondary market trading in emissions markets with bankable allowances. The MITRE Corporation, George Mason University, Case # 10-2868. https://www.mitre.org/sites/default/files/pdf/10_2878.pdf
- Park, Haeyoun. 2015. How Outsourcing Companies Are Gaming the Visa System. *The New York Times*, November 10.
- Giovanni Peri. 2012. Rationalizing U.S. Immigration Policy: Reforms for Simplicity, Fairness, and Economic Growth. Brookings Hamilton Project Discussion Paper.

Preston, Julie. 2016. Lawsuits Claim Disney Colluded to Replace U.S. Workers with Immigrants. *New York Times*, Jan 25.

Regional Greenhouse Gas Initiative, Inc. 2016. *CO₂ Allowance Auctions. Frequently Asked Questions*.
http://www.rggi.org/docs/Auctions/33/RGGI_CO2_Allowance_Auction_FAQs_July_20_2016.pdf

Rodriguez, Cristina M. 2010. Constraint through Delegation: The Case of Executive Control over Immigration Policy. *Duke Law Journal* 59:1787-1846.

Rosenberg, Mica, Stephen Nellis, and Emily Stephenson. 2017. Trump, Tech Tycoons Talk Overhaul of H1B Visas. *Reuters*, January 12. <http://www.reuters.com/article/us-trump-immigration-tech-idUSKBN14W1CZ>.

Schmalensee, R. and R. Stavins, 2013. The SO₂ Allowance Trading System: The Ironic History of a Grand Policy Experiment. *The Journal of Economic Perspectives* 27:103-122.

Simon, Julian L.. 1999. *The Economic Consequences of Immigration*, 2d ed. Ann Arbor, University of Michigan Press.

Simon, Julian L.. 1986. Auction the Right to Be an Immigrant. *New York Times*, January 28.