1 Executive Summary

Commercial health insurers pay much higher prices for health care services than public insurance programs like Medicare or Medicaid.[1] Commercial insurers pay around twice what Medicare pays for inpatient care on average, and the gap is even larger for outpatient care, as illustrated in Figure 1.1. Commercial insurers also generally pay more for physician services, although the gap is smaller.

These differences arise because commercial insurers and the Medicare program determine provider prices in different ways. In commercial insurance, provider prices are negotiated between providers and insurers. In practice, many health care providers face limited competition (e.g., Fulton 2017), which can often allow a provider to extract prices well above the minimum prices that would make serving an insurer's enrollees attractive to the provider. By contrast, Medicare generally sets provider prices administratively (i.e., via fee schedules established through legislative and regulatory processes). Historically, Medicare's prices have been set high enough to ensure that Medicare beneficiaries have a broad choice of providers (e.g., MedPAC 2020a), but policymakers' desire to contain the cost of the Medicare program has kept them well below commercial prices.

The large differential between the prices paid by Medicare and commercial insurers has led some policymakers to propose using some form of regulated or administered pricing in commercial insurance markets.[2] This paper examines three tools policymakers might use: (1) capping prices for out-of-network services; (2) regulating prices for both in-network and out-of-network services; and (3) creating a public option, a publicly operated plan that would set prices administratively and could be purchased in lieu of private plans. To gain insight into these policies, this paper develops economic models that combine economic theory with available empirical evidence. The main text summarizes the main insights from those models, and the appendices provide full mathematical details.

Figure 1.1: Average Commercial Prices as a Percentage of Medicare Prices

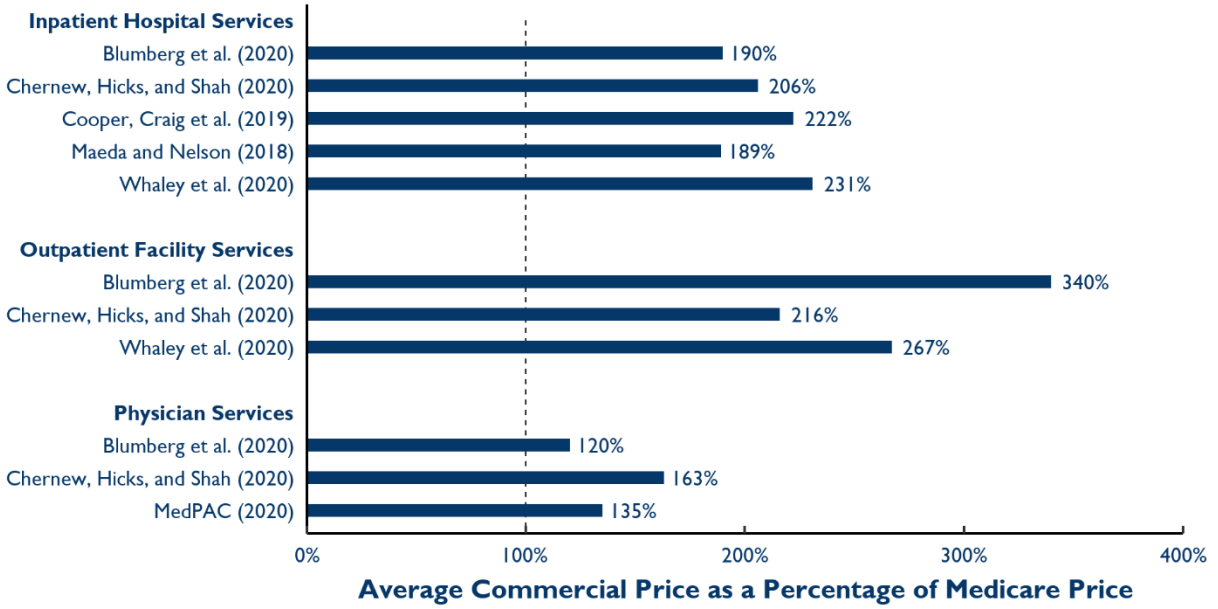

USC Schaeffer BROOKINGS

[1] Throughout, I use the term commercial insurance to encompass private insurance plans sold in the individual, small group, or large group markets, as well as self-insured group health plans offered by employers.

[2] This paper focuses on approaches to reducing the prices of health care services and largely does not consider prescription drugs in light of the major differences between prescription drugs and health care services.

The focus of this paper is understanding how these different policy tools would affect provider prices and premiums, which is of obvious interest in ongoing policy debates. Importantly, however, this paper does not seek to answer the question of whether policymakers *should* use these tools to reduce prices and, if so, how aggressively. To answer that question, it is necessary to understand how price changes caused by these policies would cause providers to change their service offerings and care delivery processes over the long run, as well as how those changes would affect the quantity and quality of the health care services patients received and the real economic resources consumed by the health care sector. Analyzing those downstream effects is beyond the scope of this paper.

1.1 Capping Prices for Out-of-Network Services

The paper first examines proposals to limit what providers can collect for out-of-network services, such as by limiting collections to some multiple of what Medicare would pay for the same services (e.g., Murray 2013; Berenson et al. 2015; Song 2017; Chernew, Pany, and Frank 2019; Melnick and Fonkych 2020b). This type of policy would directly reduce prices for out-of-network services. However, because out-of-network services account for only several percent of commercial market spending (Pelech 2020; Song et al. 2020; Chernew, Dafny, and Pany 2020), an out-of-network cap's most important effects would likely occur by changing the in-network prices negotiated by providers and insurers.

I reach the following main conclusions about this policy:

- **For services delivered in emergency situations, limiting out-of-network prices would also, in effect, limit negotiated in-network prices.** With an out-of-network cap, an insurer always has the option to break off negotiations with a provider and pay the provider the capped price. If the insurer can do this without jeopardizing its enrollees' access to the provider's services, then this option offers the insurer an attractive alternative to a negotiated agreement that would allow it to insist on an in-network price no higher than the cap. In fact, the insurer could often negotiate a price below the cap by offering the provider greater volume (via more generous coverage for the provider's services) in exchange for a lower price.

  Because federal law requires hospitals to accept patients in emergency situations, the logic above implies that an out-of-network cap could greatly reduce the prices of services delivered in emergency situations. Using the Medical Expenditure Panel Survey, I estimate that emergency department visits and ensuing inpatient stays account for 13% of health care spending for people with commercial insurance; this share is 34% for hospital services, which is arguably the service category where market power concerns are most acute.

  Naturally, the amount an out-of-network cap reduced prices would depend on where the cap was set. The gold line in Figure 1.2 illustrates the qualitative relationship between the level of the cap and the negotiated price using the formal model developed in this paper. As shown in the figure, an out-of-network cap set at a high enough level (specifically, above the provider's pre-policy charge) would have no effect on the negotiated price. But as the cap fell below that level, it would generate progressively larger reductions in the negotiated price.

- **Outside of emergency situations, an out-of-network cap may have much less scope to affect negotiated prices.** In non-emergency situations, providers are generally legally permitted to decline to treat out-of-network patients. As a result, if an insurer broke off negotiations and paid the provider the capped price, the provider could respond by turning away the insurer's enrollees (or otherwise limiting their access to its services). For this reason, an out-of-network cap would give the insurer much less leverage in non-emergency situations unless non-legal barriers (e.g., fear of public disapproval) prevented providers from turning away patients, which, as discussed further in the main text, seems unlikely.

Figure 1.2: Negotiated Prices Under an Out-of-Network Cap

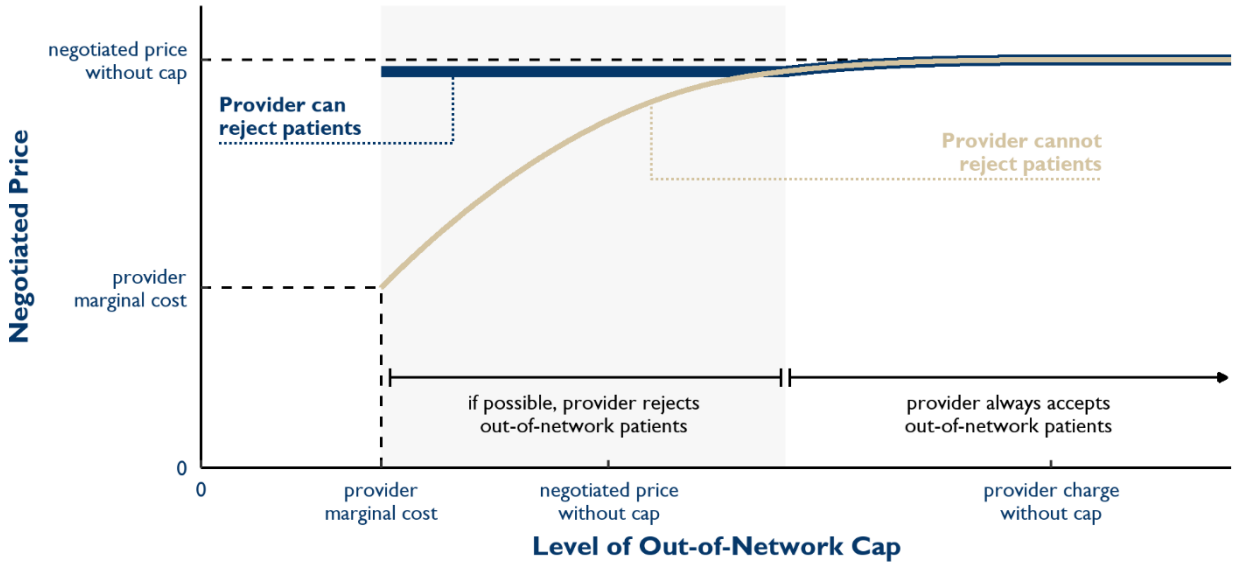

USC Schaeffer BROOKINGS

Even when a provider can turn away patients, an out-of-network cap would still weaken the provider's bargaining position to some degree. Today, a provider can generally treat *some* of an insurer's patients in the absence of a network agreement, but, with a cap, the provider's best option would often be to forgo *all* of the insurer's patients absent an agreement. How much a cap weakened the provider's bargaining position would depend on how much volume a provider can retain when out-of-network—and at what price—under the status quo. While evidence on this question is imperfect, most providers' ability to attract non-emergency out-of-network volume is likely limited.[3] Combining the fragmentary empirical evidence with the formal model developed in this paper, I conclude that an out-of-network cap could reduce negotiated prices for non-emergency services by around 10% or less.

Once again, outcomes under an out-of-network cap would depend on the level of the cap, as illustrated by the blue line in Figure 1.2. For an out-of-network cap set at a high level, serving out-of-network patients would remain lucrative, so the provider's best option would be to accept out-of-network patients at the capped price, and incrementally tightening the cap would cause small reductions in the negotiated price. But for a low enough cap, it would be in the provider's interest to forgo out-of-network patients, and further tightening the cap would have no effect on negotiated prices or other outcomes of interest.

- **Capping out-of-network prices could make it harder to obtain non-emergency out-of-network services.** The analysis above concludes that, for a stringent enough out-of-network cap, providers would wish to turn away out-of-network patients in order to protect their bargaining leverage vis-à-vis insurers. In practice, providers might find ways to accept out-of-network patients in cases where doing so would not undermine their bargaining

[3] One notable exception is services delivered by ancillary physicians (radiologists, anesthesiologists, pathologists, or assistant surgeons) during a hospitalization. As noted in recent debates over surprise billing, these physicians often retain substantial volume even when out of network (e.g., Adler, Fiedler, Ginsburg, Hall, et al. 2019; Cooper, Scott Morton, and Shekita 2020). Spending on these services is a modest, but not trivial, share of commercial spending (Cooper et al. 2020; Duffy et al. 2020).

position (e.g., uninsured patients and traveling patients). Nevertheless, it still might become harder for insured patients to routinely access care from out-of-network providers.

The paper also briefly considers a related policy that would place *both* a cap on what providers can collect for out-of-network services *and* a floor on what insurers must pay for out-of-network services (and how much coverage insurers must offer for that care). Unlike an out-of-network cap, this policy has the potential to *increase* negotiated prices if the floor is set at a high level. In particular, a provider has no reason to accept a negotiated price below the floor price because, even if negotiations break down, the policy's floor on how much coverage the insurer must offer for out-of-network care would ensure that the provider could continue to attract significant volume and be paid the floor price.

Notably, in non-emergency situations, this type of policy could increase prices on average even if the floor is set at a moderate level, such as the average negotiated price under the status quo. This is because the floor portion of the policy would place upward pressure on the prices negotiated by low-priced providers, but high-priced providers could keep the cap portion of the policy from substantially reducing the prices they receive by threatening to turn away out-of-network patients.

1.2 Regulating Both In-Network and Out-of-Network Prices

Because an out-of-network cap would likely have little effect on negotiated prices for non-emergency services, policymakers might wish to consider policies that would directly regulate both in-network and out-of-network prices. The next section of the paper thus considers two approaches to doing so, which I call the "comprehensive price cap" and "default contract" approaches.

A comprehensive price cap, as I define it here, would directly limit the amounts providers can receive for delivering health care services, both in and out of network (e.g., Skinner, Fisher, and Weinstein 2014; Murray and Berenson 2015; Blumberg et al. 2019; Roy 2019; Chernew, Dafny, and Pany 2020). The paper reaches the following conclusions about the effects of a comprehensive price cap:

- **A comprehensive price cap could reduce prices for all health care services, including in settings where providers can turn away out-of-network patients.** When providers must accept out-of-network patients, a comprehensive price cap and an out-of-network cap would be equally effective in reducing prices. But when providers can turn away out-of-network patients, a comprehensive price cap would have much greater scope to affect negotiated prices than an out-of-network cap. While a provider could keep the out-of-network portion of a comprehensive price cap from undermining its bargaining position by threatening to turn away out-of-network patients, the in-network portion of the cap would prevent the provider from translating a strong bargaining position into high prices.

- **Under a comprehensive cap, providers could use the leverage that they could not translate into higher prices to extract other concessions, which could undermine the cap or have other undesirable effects.** As noted above, a comprehensive price cap would reduce prices partly by directly limiting the prices providers and insurers could agree to rather than by reducing how much leverage providers held in network negotiations. But the leverage that providers could not translate into higher prices would not disappear, and providers could use that "excess" leverage to extract other types of concessions from insurers.

  Providers might, for example, use their excess leverage to resist contract provisions intended to discourage inefficient utilization, such as prior authorization requirements or new payment models. Providers' incentives to increase volume was historically a concern under state hospital rate setting systems (e.g., Pauly and Town 2012; Murray and Berenson 2015).

Alternatively, providers might circumvent the cap by demanding insurers pay higher prices for service lines where the cap does not apply (or does not bind). For example, a health system with a high-priced flagship hospital could seek higher prices for its lower-priced community hospitals or its physician practices; systems might also accelerate acquisitions of hospitals or physician practices to maximize their ability to use this strategy. Evasion concerns would likely also require policymakers to limit use of alternative payment models, like bundled payments or shared savings contracts, since such contracts could be used to "hide" payments to providers. Policymakers would have options for addressing these problems, but it is unclear how effective they would be, and some might have undesirable side-effects of their own.

Motivated by the enforcement challenges that could arise under a comprehensive price cap, this paper also considers an alternative way of regulating health care prices that I call the "default contract" approach.[4] Under this approach, the government would publish a model network agreement (the "default contract") that specified both the prices the insurer would pay the provider *and* a minimum level of access the provider would be required to offer to the insurer's enrollees. A provider would be required to enter a default contract with any insurer that requested one, but providers and insurers would also be allowed to negotiate any alternative payment terms they wished. An insurer would be permitted to request a default contract with some providers but not others at its discretion.

The paper reaches the following main conclusions about the default contract approach:

- **A default contract approach could reduce prices for all health care services, while avoiding the main enforcement challenges of a comprehensive price cap.** Under a default contract policy, the insurer would always have the option to break off negotiations and give its enrollees access to the provider's services via a default contract. This option would allow the insurer to insist on prices at or below those in the default contract, at least if the default contract's access standards were reasonably stringent and effectively enforced.

  Importantly, the default contract approach would limit prices by directly weakening a provider's leverage in network negotiations, rather than by limiting the provider's ability to translate leverage into high prices. For this reason, unlike a comprehensive price cap, it would not spur provider efforts to use leverage they cannot translate into higher prices to extract other concessions, such as higher volume or higher prices for service lines not subject to the price cap. Nor would it be necessary to limit use of alternative payment models.

- **The core challenge of the default contract approach would be enforcing the access standards.** The default contract approach would only be effective in reducing prices if implementing a default contract gave an insurer's enrollees real access to the provider's services, which would require the default contract's access standards to be effectively enforced. While not easy, enforcing these access standards would likely be easier than overcoming the various enforcement challenges that could arise under a comprehensive price cap.

  Notably, enforcement efforts could focus on a single problematic behavior—provider attempts to turn away enrollees covered under a default contract—rather than the many different problematic behaviors that could arise under a comprehensive price cap. Additionally, provider compliance with the access standards would be comparatively straightforward to monitor directly via insurer or consumer complaints and, if necessary, audit studies.

  Importantly, a default contract policy could be effective in reducing prices even if access standards were enforced imperfectly. While imperfect enforcement would reduce the leverage

[4] Glied and Altman (2017) describe a version of this approach that would apply to a narrow subset of hospital services.

insurers derived from the ability to demand a default contract, policymakers could compensate for imperfect enforcement to some degree by specifying lower prices in the default contract.

1.3 Creating a Public Option

Another way to introduce regulated or administered pricing in the commercial market is to create a "public option," a publicly operated plan that consumers could purchase in lieu of a private plan. Introducing a public option was considered during the debate over the Affordable Care Act (ACA), and President-elect Biden's campaign platform included a public option. Many Congressional and think tank proposals also envision introducing a public option (T. Neuman et al. 2019).

Public option proposals vary widely in design. This paper focuses on a public option that would pay health care providers some percentage of the prices Medicare pays providers, require providers to accept public option patients, and charge a premium that covers its average costs. However, I also discuss how the effects of alternative public option designs might differ.

Market outcomes with a public option, including the prices providers received, the premiums enrollees paid, and the market share captured by the public option would depend on how private plans—and, particularly, private plans' negotiations with providers—changed in response to creation of a public option. To gain insight on these dynamics, this paper develops a formal model of health insurance markets in the presence of a public option. The main text presents the main insights from that model and the results of simulations using that model. Appendix B presents full details.

This analysis reaches the following main conclusions about the effects of introducing a public option:

- **If a public option was much more attractive to consumers than existing private plans, then private plans would end up paying providers prices close to the public option's prices.** The introduction of a public option that was much more attractive to consumers than existing private plans would reshape provider-insurer negotiations in two important ways. First, consumers would be unwilling to pay too much more for a private plan than for the public option, which would force private plans to set premiums close to the public option's premium. That, in turn, would make it unprofitable for private plans to pay providers prices too far above the public option's, making insurers willing to walk away from network negotiations rather than pay prices that high. Second, providers would recognize that if they did not join private plans' networks, some of their patients would instead enroll in the public option, and they would be paid the public option's prices. Thus, providers would be willing to walk away from negotiations rather than accept prices too far below the public option's prices.

  Virtually any coherent economic model of provider-insurer bargaining predicts that a provider and insurer will negotiate a price that lies between the maximum price that makes an agreement profitable for the insurer and the minimum price that makes an agreement profitable for the provider. Thus, the considerations above imply that prices providers and insurers negotiated would end up neither too far above nor too far below the public option's prices. This conclusion contradicts assumptions in some prior analyses that introducing a public option would not meaningfully change the prices private plans could negotiate (Antos and Capretta 2019; FTI Consulting 2019; Koenig et al. 2019; Schaefer and Moffit 2020). This analysis also suggests that a public option that paid most providers less than existing private plans but paid some providers more could increase the prices that private plans paid those specific providers, even as it reduced the prices that private plans paid providers overall.

  Importantly, the conclusions above depend on the public option being a strong competitor for private plans. If a public option had non-price cost disadvantages relative to private plans that partially offset its pricing advantages (a possibility discussed below), then it would set

correspondingly higher premiums and do less to constrain the premiums private plans set and the prices they paid providers. Indeed, if the public option had non-price cost disadvantages large enough to fully offset its pricing advantages, it would likely attract little enrollment and have little effect on market outcomes. Similarly, a public option that paid all providers more than existing private plans would also have little effect on market outcomes.

- **A public option that paid providers less than existing private plans could both offer consumers a new lower-premium option and reduce the premiums of private plans.** The preceding discussion implies that, in cases where the public option was more attractive to consumers than existing private plans, both private plans and the public option would pay providers prices that were reasonably close to the public option's prices. Thus, if a public option paid providers less than existing private plans and did not have large non-price cost disadvantages, the premiums set by both the public option and the private plans competing with it would likely be lower than the premiums of existing private plans.

  Notably, employer plans pay providers very different prices in different parts of the country (e.g., Chernew, Hicks, and Shah 2020), and it is generally believed that individual market plans pay providers less than employer plans (e.g., Blumberg et al. 2020). Thus, if a public option paid providers the same prices in all settings, it would likely have different effects on premiums and prices in different geographic areas and different insurance markets. Specifically, it would tend to reduce premiums the most in areas and markets where private plans currently pay the highest prices, but generate smaller, if any, savings in lower-priced areas or markets.

- **A public option would differ from private plans in ways other than what it paid providers, including non-price determinants of plan costs (e.g., utilization, non-claims costs, risk selection, and diagnosis coding) and how it set premiums.** In particular, experience from Medicare Advantage (e.g., Curto et al. 2019) suggests that a public option would have higher utilization than its private competitors for comparable enrollees, at least in the individual market, where private plans are often tightly managed. On the other hand, data on non-claims expenses in traditional Medicare and existing private plans suggests that a public option might have lower non-claims expenses than competing private plans.

  In the individual and small group markets, the public option would likely also differ from private plans in what types of enrollees it attracted and how aggressively it coded diagnoses for risk adjustment purposes. Experience from Medicare Advantage suggests that private plans might attract a healthier mix of enrollees and succeed in making comparable enrollees look sicker for risk adjustment purposes (e.g., Curto et al. 2019; Geruso and Layton 2020). Both factors would tend to increase the public option's costs relative to private plans.

  A public option would also set premiums differently from private plans. While the public option would set a premium to cover its average costs, private plans set premiums to maximize their profits. Correspondingly, private plans would set premiums that incorporate a markup over their costs, ceding some enrollment to the public option in exchange for positive margins. Because of this difference in premium-setting behavior, introducing a public option could be particularly consequential in areas with few competing insurers (and, thus, high markups).

Figure 1.3: Public Option Market Share

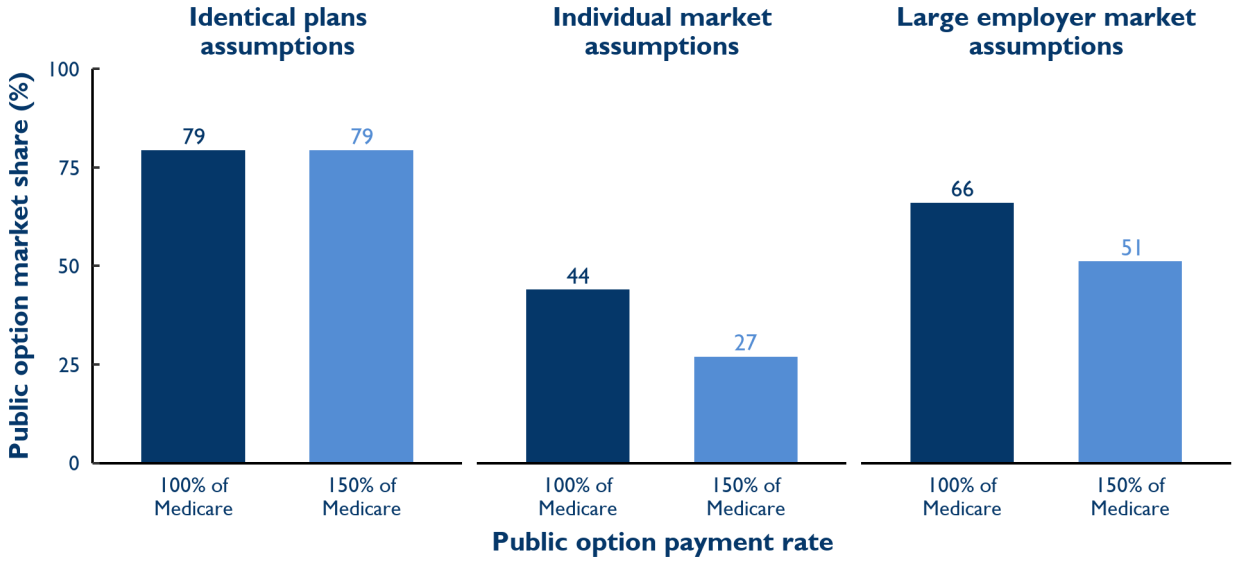

Note: The figure reflects a scenario in which existing plans pay providers 180% of Medicare rates, on average.

USC Schaeffer BROOKINGS

- **The public option's market share could vary widely depending on how it compared to private plans, but private plans would retain substantial enrollment in most plausible scenarios.** Figure 1.3 illustrates this fact using the simulations conducted in this paper. The figure examines several scenarios in which a public option is introduced in a market where existing private plans pay providers 180% of Medicare rates. (As discussed in the main text, the model used here includes only a single private plan, which may cause it to overstate private plans' premiums and understate their market share. However, while the results displayed in Figure 1.3 should not be taken too literally, they do help illustrate how and why the public option's market share would likely vary across different scenarios.)

  The first set of scenarios assumes that the public option and private plan are identical, except for determining provider prices and premiums differently. These scenarios are unrealistic but offer a useful benchmark. In these scenarios, the public option captures about four-fifths of the market. This occurs because the private plan's premium incorporates a markup over its costs and thus charges a higher premium despite having an identical cost structure.

  The second set of scenarios reflects assumptions plausible for a public option offered in the individual market. For these scenarios, I assume that the public option has higher utilization than the private plan, attracts sicker enrollees, and codes diagnoses less aggressively in risk adjustment, but has lower non-claims costs. The private plan also has a narrower network that allows it to negotiate prices modestly below the public option's prices. Thus, the public option charges a higher premium than the private plan and captures only a minority of the market.

  The final scenarios reflect assumptions plausible for a public option offered to large employers. Since private plans in the employer market have broader networks and weaker utilization controls, I assume that a public option offered to employers would have a smaller utilization disadvantage; the private plan's broader network also causes it to pay prices closer to the public option's. Additionally, because I assume that an employer market public option would only offer third-party administrator services, risk selection is no longer relevant. Consequently, the

public option's premium is lower relative to the private plan's than in the second set of scenarios, so the public option captures half or more of the market.

The consequences of introducing a public option would differ if the design of a public option differed from the one considered in this paper's primary analyses. In particular:

- **If providers could opt out of serving public option patients, many providers might do so, potentially leading to a very different market equilibrium.** A provider that opted out of the public option would become more valuable to private plans—because private plans could now offer exclusive access to the provider's services—and thus be able to negotiate higher prices with private plans. While opting out would also have costs for providers, primarily lost profits on public option volume, it is plausible that many providers would opt out, at least if the public option set low payment rates. Providers that command high prices under the status quo would likely have the most to gain by opting out.

  The consequences of provider opt outs would depend on how widespread they were. If the public option's network ended up far narrower than existing private plans, then introducing a public option might have little effect on market outcomes, either because the public option would attract little enrollment or because policymakers would be forced to pull the public option from the market. If the public option ended up with a narrow, but viable network, the situation is more complex. Relative to the case where providers must participate in the public option, private plans would likely pay providers more and charge higher premiums, while the public option might have lower utilization and suffer from less adverse selection (Liu et al. 2020). The net change in the public option's market share relative to the case with mandatory provider participation from a narrower network and a lower relative premium is uncertain.

- **If a public option negotiated prices with providers rather than setting them administratively, it is doubtful that a public option would pay lower prices than existing private plans.** If policymakers wished to make participation in the public option voluntary for providers but still allow the public option to attract a broad network, they could implement a public option that set prices through negotiation with providers, rather than administratively. However, there is little reason to believe that a public option would be able to negotiate lower prices than existing private plans since it would not have access to any negotiating tools beyond those available to private plans. A public option might still charge a modestly lower premium by virtue of setting a premium that does not incorporate a profit margin or by having lower non-claims costs, but these advantages might be offset in practice if the public option had higher utilization or experienced adverse selection.

1.4 Effects on Provider Networks

Most of the analysis in this paper focuses on how a price cap or public option might affect prices and premiums generally, but these policies would likely have different effects on plans with broader and narrower networks and change what types of networks enrollees select. While this paper does not offer a full analysis of potential effects of these policies on plans' networks, I reach the following qualitative conclusions about the effects of a price cap or a public option on these outcomes:

- **These policies would likely reduce the difference in premiums between broad and narrow network private plans.** Because all of the policies considered in this paper would reduce the overall level of provider prices, they would reduce the savings insurers could realize using narrow networks; when the overall level of prices is lower, an insurer's scope to use a narrow network to negotiate still lower prices is smaller, and the potential savings from using a narrow network to steer enrollees away from high-priced or high-utilizing providers is

smaller too. Correspondingly, these policies are likely to shrink the gap in premiums between broad and narrow network plans. For a public option offered in the individual market or small group market, changes in risk selection patterns could also affect the relative premiums of broad and narrow network plans, although the direction of this effect is uncertain.

- **While reductions in the relative premiums of broad network plans would generally push consumers toward broader networks, some factors could push in the opposite direction.** In particular, the price cap policies would reduce consumers' exposure to balance billing when they receive out-of-network care, which could make narrow network plans modestly more appealing, perhaps partially offsetting the fact that opting for a narrow network plan would now offer smaller premium savings. Under a public option, the public option might siphon off many enrollees who prefer broad network plans, which could cause *private plan* enrollment to shift toward narrower networks even though narrow networks would now offer smaller premium savings, although *overall* enrollment (inclusive of public option enrollment) would still likely shift toward broader networks.

Any shift toward broader networks in private plans would tend to partially offset the downward pressure on average provider prices and premiums created by a price cap or public option.

1.5 Strategies for Ensuring Provider Compliance

Either a price cap or a public option would impose requirements on health care providers, and a natural question is how those requirements would be enforced. Policymakers would have two broad categories of options. First—and most straightforward—they could directly penalize non-compliant providers. For example, policymakers could fine non-compliant providers, and state policymakers could consider making compliance a condition of provider licensure.

Second, federal policymakers could require providers to comply with a price cap or accept patients under a public option in order to serve patients with various forms of federally subsidized coverage. A narrow version of this approach might encompass only public programs like Medicare and Medicaid, while a broader version could also encompass private insurance plans offered on the group and individual markets, which are subsidized via the tax exclusion for employer-provided coverage and the ACA's Marketplace subsidies. Naturally, the more types of subsidized coverage included, the more successful this approach would likely be in ensuring compliance with the price cap or public option.

Importantly, one risk of this approach is that providers might opt out of the relevant forms of publicly subsidized coverage rather than comply with the price cap or public option. That concern would be most acute for a price cap that was set at a low level or that affected a broad array of services, as well as public option that paid low prices. It would also be larger for a price cap or a public option that was implemented in the group market in addition to the individual market. On the other hand, it would tend to be smaller if all (or almost all) forms of federally subsidized coverage were included in this type of approach. Virtually all existing coverage is federally subsidized in some way, so being locked out of all forms of federally subsidized coverage would likely be viable for few, if any, providers.

1.6 Experience from Medicare Advantage

Experience with most of the policy tools considered in this paper is relatively limited in the United States. But the Medicare program is an important exception. In Medicare, private Medicare Advantage (MA) plans compete alongside traditional Medicare, which plays the role of a public option, and providers are subject to an out-of-network cap set at traditional Medicare rates when treating MA enrollees. The Medicare program thus offers an interesting empirical setting in which to assess and apply the largely theoretical analysis presented in the rest of this paper.

A striking fact is that MA plans pay hospitals and physicians prices very close to traditional Medicare's prices in almost all cases, a stark contrast with the much higher and widely varying prices paid by commercial plans (Berenson et al. 2015; Baker et al. 2016; Trish et al. 2017; Maeda and Nelson 2018; Pelech 2020). Applying this paper's theoretical analysis to MA yields two conclusions, which offer both some support for this paper's analysis and some insight into dynamics in MA:

- **The presence of traditional Medicare can largely explain the prices observed in MA.** Medicare program rules make it impossible for institutional providers to turn away traditional Medicare patients while still serving MA enrollees, and traditional Medicare's large market share likely makes turning away traditional Medicare patients unattractive to physicians too. Traditional Medicare is thus analogous to a public option with mandatory provider participation. Correspondingly, the analysis of a public option in this paper implies that the presence of traditional Medicare should allow MA plans to negotiate prices close to traditional Medicare's, consistent with the prices actually observed in MA. This echoes some prior analyses of MA that have also posited a major role for traditional Medicare in explaining the prices observed in MA (e.g., Berenson et al. 2015; Trish et al. 2017).

- **While the MA out-of-network cap likely plays at least a supporting role in explaining the prices observed in MA, that role may be smaller than sometimes suggested.** There do not appear to be clear legal or other barriers keeping providers from turning away out-of-network MA enrollees (or otherwise limiting access) in non-emergency situations. Thus, the analysis of an out-of-network cap in this paper suggests that the out-of-network cap likely has only modest effects on the prices MA plans can negotiate for non-emergency services. This conclusion differs from prior work that assigns the out-of-network cap a more central role in shaping prices in MA (e.g., Maeda and Nelson 2018; Pelech 2020).

  The presence of an out-of-network cap may nevertheless play a supporting role in shaping negotiated prices in MA. Even when providers can turn away patients, an out-of-network cap does have some limited scope to reduce prices. This may matter in cases where competitive pressure from traditional Medicare leaves the prices negotiated by MA plans modestly above traditional Medicare's prices. In these cases, the out-of-network cap may push negotiated prices the rest of the way toward traditional Medicare's, which may help explain why MA prices are *uniformly* close to traditional Medicare's across different providers and geographic areas.

1.7 Conclusion

The analysis in this paper demonstrates that an appropriately designed price cap or public option can reduce the prices of health care services. It also offers some guidance on how policymakers that wished to use a price cap or a public option to reduce prices should choose among these policies:

- **Neither an out-of-network cap nor a comprehensive price cap is likely to be policymakers' best option to reduce provider prices.** It is questionable at best whether an out-of-network cap could reduce the prices of services delivered in non-emergency situations, and it could reduce patients' ability to access out-of-network care. A comprehensive price cap could, on paper, reduce prices in all settings, but enforcement challenges might threaten the integrity of the cap, and the cap could have various undesirable side effects, including increased utilization, greater consolidation, and less adoption of alternative payment models. By contrast, a default contract policy could reduce prices in all settings while avoiding the main enforcement challenges and undesirable side-effects of the other approaches.

- **If policymakers' sole goal is to reduce provider prices, a default contract policy is a simpler and more flexible tool than a public option.** A default contract policy could

be applied solely to specific services (à la Glied and Altman 2017; Roy 2019), whereas a public option would need to set prices for all types of services and, correspondingly, would affect prices for all types of services. Additionally, a default contract policy could be targeted primarily at the highest-priced providers (à la Chernew, Dafny, and Pany 2020) by specifying high prices in the default contract. By contrast, a public option that paid all providers more than existing private plans would be uncompetitive and thus have little or no effect on prices, and a public option that paid somewhat lower prices would increase prices received by low-priced providers in addition to reducing the prices received by high-priced providers.

Implementing a public option would also entail operating an insurance plan, which would be administratively complex. Related, if the public option had disadvantages in utilization, risk selection, or diagnosis coding, that could undermine its ability to reduce provider prices.

- **A public option can address insurer market power, which the price cap policies cannot.** Policymakers may have goals other than reducing provider prices. Notably, many insurance markets are concentrated (Fulton 2017), which can allow insurers to charge higher premiums (e.g., Dafny, Duggan, and Ramanarayanan 2012; Dafny, Gruber, and Ody 2015). Introducing a public option could reduce premiums by forcing insurers to accept smaller profit margins (although such margins are already generally modest) or by creating pressure for insurers to operate more efficiently along other dimensions. These considerations can provide a rationale for implementing a public option instead of or in addition to a price cap.

While this paper focuses on the substantive effects of these policies, policymakers would also need to consider the political feasibility of alternative policy approaches. Political considerations might be particularly important to the choice between a public option and some form of price cap. Notably, introducing a public option would threaten the interests of health insurers in addition to health care providers and thus could spark broader industry opposition. However, health insurers are deeply distrusted by the public (Commonwealth Fund, New York Times, and Harvard T.H. Chan School of Public Health 2019; KFF 2020a), so a public option that offered consumers a concrete alternative to private insurance plans could have broad public appeal. Indeed, this antipathy for private insurers may well be part of the reason that public opinion survey data show that public option proposals command broad public support (Kirzinger, Kearney, and Brodie 2020).