INSEAD The Business School for the World®

Working Paper

2021/31/TOM

(Revised version of 2020/44/TOM)

Be the Match: Optimizing Capacity Allocation for Allogeneic Stem Cell Transplantation

Sundara Natarajan, Panchanatham
INSEAD, sundaranatarajan.panchanatham@insead.edu

Michael Freeman
INSEAD, michael.freeman@insead.edu

Harry Groenevelt
University of Rochester, groenevelt@simon.rochester.edu

Sameer Hasija
INSEAD, sameer.hasija@insead.edu

Problem definition: Treating many blood-related diseases requires transplantation of genetically compatible hematopoietic stem cells (HSCs) extracted from the bone marrow (BM) of live donors or the umbilical cord blood (CB) of babies. To facilitate the search for HSCs, institutions known as BM registries collect the details of potential donors and CB banks store units of CB. This paper focuses on the problem of joint optimization of the capacity of these two institutions.
Academic: With over 10 million genetic variants, limited inventory relative to this variety, and random replenishment, BM registry and CB bank compositions are random, inter-dependent, and change nondeterministically over time. Furthermore, BM and CB differ in their supply, costs, genetic matching criteria, and influences on medical outcomes, giving rise to important trade-offs such that neither is preferred exclusively to the other. Jointly determining the optimal capacity of both sources is therefore both technically challenging and has immediate policy implications.
Methodology: We develop a simulation-based approach to estimate the temporal variation in matching probabilities before incorporating the associated regression parameters into a mathematical model closely matching the research context. Results are contrasted against a simplified mathematical model, highlighting the importance of the dynamic setup.
Results: Inventories of 17.5 million registered BM donors and 335 thousand CB units are estimated as optimal for the U.S. population under reasonable assumptions. Expanding capacity to these levels would satisfy 33% of the currently unmet demand, increasing the transplantation rate to 98.7% and delivering $770 million of extra social surplus annually.
Managerial Implications: Rigorous policy analyses are imperative for designing evidence-based, cost-effective policies that deliver societal benefits. To this end, we provide quantitative evidence in support of calls for further expansion of the national BM registry and CB banks in the U.S. We also propose annual recruitment targets for BM donors and CB units to maintain the two institutions at their suggested levels.

Keywords: Healthcare; Allogeneic Transplanation; Stem Cells; Bone Marrow; Cord Blood; Optimization; Simulation
History: Last updated on June 15, 2021.

Be the Match: Optimizing Capacity Allocation for Allogeneic Stem Cell Transplantation

Sundara Natarajan, Panchanatham
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, sundaranatarajan.panchanatham@insead.edu,

Michael Freeman
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, michael.freeman@insead.edu,

Harry Groenevelt
Simon Business School, University of Rochester, NY 14627, groenevelt@simon.rochester.edu,

Sameer Hasija
INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, sameer.hasija@insead.edu,

Problem definition: Treating many blood-related diseases requires transplantation of genetically compatible hematopoietic stem cells (HSCs) extracted from the bone marrow (BM) of live donors or the umbilical cord blood (CB) of babies. To facilitate the search for HSCs, institutions known as BM registries collect the details of potential donors and CB banks store units of CB. This paper focuses on the problem of joint optimization of the capacity of these two institutions.

Academic: With over 10 million genetic variants, limited inventory relative to this variety, and random replenishment, BM registry and CB bank compositions are random, inter-dependent, and change nondeterministically over time. Furthermore, BM and CB differ in their supply, costs, genetic matching criteria, and influences on medical outcomes, giving rise to important trade-offs such that neither is preferred exclusively to the other. Jointly determining the optimal capacity of both sources is therefore both technically challenging and has immediate policy implications.

Methodology: We develop a simulation-based approach to estimate the temporal variation in matching probabilities before incorporating the associated regression parameters into a mathematical model closely matching the research context. Results are contrasted against a simplified mathematical model, highlighting the importance of the dynamic setup.

Results: Inventories of 17.5 million registered BM donors and 335 thousand CB units are estimated as optimal for the U.S. population under reasonable assumptions. Expanding capacity to these levels would satisfy 33% of the currently unmet demand, increasing the transplantation rate to 98.7% and delivering $770 million of extra social surplus annually.

Managerial Implications: Rigorous policy analyses are imperative for designing evidence-based, cost-effective policies that deliver societal benefits. To this end, we provide quantitative evidence in support of calls for further expansion of the national BM registry and CB banks in the U.S. We also propose annual recruitment targets for BM donors and CB units to maintain the two institutions at their suggested levels.

Key words: Healthcare; Allogeneic Transplanation; Stem Cells; Bone Marrow; Cord Blood; Optimization; Simulation
History: Last updated on June 15, 2021.

1. Introduction

The UN, in their sustainable development goals (SDG), call for decisive action to promote healthy lives and wellbeing for all, including a call to cut premature mortality from non-communicable diseases drastically. Providing citizens with cost-effective access to life-saving technology and treatment is essential to attaining that goal. Yet every year, hundreds of thousands of people worldwide are diagnosed with hemato-

logical (blood-related) diseases such as leukemia and lymphoma (types of cancer), severe aplastic anemia (failure of bone marrow functioning), inherited immune system disorders, and inherited metabolic disorders. For many patients with such diseases, transplanted blood-forming stem cells, known as hematopoietic stem cells (HSCs), from healthy donors can improve their likelihood of survival. This process requires collecting HSCs from a donor whose immune system is compatible with that of the recipient and then transplanting these HSCs into the patient in a procedure known as an allogeneic stem cell transplantation.

Generally, there are three main sources of HSCs for allogeneic transplants: 1) related bone marrow (BM) donors, i.e., a close relative of the patient, usually a sibling; 2) matched unrelated bone marrow donors (MUDs), i.e., donors from the general population whose genes match those of the patient; and 3) umbilical cord blood (UCB), i.e., stem cells extracted from the umbilical cords of newborn babies whose genes have a close match with the patient's. Empirically it has been observed that 70% of patients do not find a related donor and hence their prognosis depends on finding an unrelated BM donor or compatible UCB cells held in storage (Gragert et al. 2014). Based on 2019 population demographics, nearly 13,000 patients are therefore estimated to require allogeneic transplants from unrelated donors every year in the U.S. alone (Health Resources and Services Administration 2020b).

To facilitate the search for unrelated HSCs, different countries have established institutions known as bone marrow registries (BMRs) and cord blood banks (CBBs). BMRs store the contact and genetic information of voluntary BM donors, while CBBs store previously collected cord blood units (CBUs) in cryogenic facilities. In the U.S., for example, the National Marrow Donor Program (NMDP) currently maintains a registry of almost nine million potential BM donors. The contact details of these potential donors are made available to transplant centers through a program called "Be The Match." In addition, approximately 265,000 CBUs are stored in public CBBs in the U.S., most of which are listed under Be The Match and made available on request.[1] However, despite the large size of these institutions relative to the total demand, Be The Match reports that a patient's likelihood of finding a matching donor in the registry ranges only from 23% to 77% depending on ethnic background (Be The Match 2020b).

One reason put forward for the low matching probability is that, at present, the BMRs and CBBs contain an insufficient number of donors and CBUs (Health Resources and Services Administration 2020c). However, increasing the size of these institutions is not a trivial decision as doing so would impose a significant cost burden on the healthcare system. Moreover, the benefit of increasing the BMR and CBB size, in terms of the additional lives saved, is not straightforward to determine due to the complexities associated with finding a compatible match for a recipient. A significant body of literature in the fields of medicine and economics (discussed in detail in Section 3) has therefore focused on demonstrating how the matching

[1] Current BM donor and CBU numbers correct as of 23 August 2020 (World Marrow Donor Association 2020).

probabilities differ depending on the size of the BMR or CBB or on estimating target annual recruitment rates. To date, however, any attempt made at optimizing the size of the two institutions has always considered each in isolation but never both together. Yet there are important trade-offs between expanding the BMR and expanding the CBB, making it crucial from a policy and practical perspective to consider the two alternative sources in unison. This paper looks to fill this gap in the literature.

In particular, the most important practical difference between stem cells from BM and those from UCB is that successful transplantation requires different levels of donor-recipient genetic compatibility (medically referred to as histocompatibility). The specific biological details that describe a so-called "match" are given in Section 2, but essentially this is determined by the number of matches between six alleles that make up the phenotypes of the patient and donor. For a BM transplant to be successful, a perfect match (i.e., a 6/6 phenotype match) is required. In contrast, one advantage of UCB cells over BM is that matches as low as 4/6 can still lead to a successful transplant (though a perfect match is still preferred to 5/6 and 4/6 matches).

In addition to their matching requirements, these two sources of HSCs also differ in their costs (e.g., of recruitment, extraction, and storage), availability, and, importantly, medical effectiveness. (More details are provided in Section 2, Section 4, and Section EC.3 of the e-companion.) It is the latter, in particular, that usually makes BM cells a physician's first choice. That said, for patients with rarer phenotypes it may be impossible to find the required 6/6 match in the BMR despite the large size of the registry. In such cases, using stem cells from the CBB may be the only viable option. For this reason, there is significant tension between expanding the size of the BMR to increase the chance of a rare match versus operating an expanded CBB containing CBUs that are more flexible in catering to these rare cases.

From a public health and policy perspective, then, the optimal sizing of these two institutions, each different in their operational aspects and medical effectiveness, represents an important and challenging decision. However, as we discuss later, solving this dual optimization problem is complicated by the fact that the inventory of CBUs in the CBB changes dynamically over time in ways that are difficult to predict. The specificity of the problem thus necessitates a practical approach that calls for close alignment between the context and the research procedure, and which requires viewing the problem through an operations lens to provide practically relevant findings with immediate implications for public health policy.

The key contribution of this paper therefore lies in finding the optimal size of each institution given a fixed annual demand for allogeneic transplants. First, we mathematically model the net annual benefits arising from transplants from the BMR alone and optimize the BMR size for the U.S. population. This builds on existing work (e.g., by Bergstrom et al. 2009) by incorporating losses due to annual attrition and no-shows that have been previously overlooked. Thereby, we also identify the annual rate of donor recruitment needed to compensate for donor loss. Second, we formulate an optimization model that finds the optimal BMR and

CBB sizes (under an extreme policy in which all CBUs are discarded at the end of each year) when both sources are used together to cater to HSC demand. Finally, we develop and analyze a simulation model to study the temporal dynamics of the system with both sources of supply under real-world conditions. A regression model is then fitted to the output of the simulation and achieves near-perfect predictive accuracy. The predictive equation is then incorporated into a mathematical model that is optimized to provide long-term solutions to the problem at hand.

Overall, we find that the optimal BMR in the U.S. will comprise approximately (approx.) 17.5 million registered donors, who are able to satisfy approx. 92.8% of the total annual demand for HSCs with an annual operating cost of approx. $73.5 million USD. For the remaining patients who are unable to find a match in the BMR, we estimate a total public CBB inventory of approx. 335 thousand CBUs to be sufficient. These CBUs can satisfy approx. 5.9% of the total annual demand (or, equivalently, 82.3% of the spillover from the BMR) with an annual operating cost of approx. $52 million USD.

From a policy perspective, despite recent suggestions that the number of CBUs currently in storage in public banks in the U.S. is too high, we find that an approx. 26% expansion in CBB capacity is still required before the inventory level approaches optimal. Meanwhile, the current total number of registered BM donors in the U.S. (estimated at approx. 11.5 million donors[2]) would need to increase by 52% to reach the optimal value. Our results therefore suggest that approx. $629 million in additional funding should be made available for recruiting new donors, with a further $141 million allocated for the collection of new CBUs. These up-front costs, although high, are justified by the increase in the number of transplants per year, with 33.0% of the previously unmet demand satisfied under the recommended policy. These transplants deliver approx. $770 million in additional value annually, which, after accounting for *all* costs, translates to an increase of approx. $6.9 billion in total social surplus over a 10-year time horizon.

From a practical perspective, our findings can also guide targets for annual donor recruitment and UCB collection. With a 4% BM donor attrition rate per year and a 20-year expiry for CBUs, we estimate that approx. 700 thousand new BM donors must be recruited each year and 35 thousand new CBUs collected to maintain the respective institutions at their recommended sizes. Perhaps more importantly, we contribute by providing an approach that can be easily adapted so that these targets can be refined over time as new or more accurate data becomes available (e.g., on the distribution of phenotypes in the population or other parameters used as inputs). Finally, we note that while our recommendations are specific to the U.S. context, our estimation routine can be readily applied to healthcare systems outside of the U.S. to guide capacity and recruitment decisions.

[2] Author's calculations based on the number of registered donors in the U.S. across the four largest BMRs: Gift of Life (≈375,000 – World Marrow Donor Association 2020), NMDP (≈8.95 million – World Marrow Donor Association 2020), Salute to Life (close to 1 million – Be The Match 2020a), and DKMS (≈1.13 million – DKMS 2020).

2. Background Information: Biology and Context

In this section, we provide more information on the genetic compatibility behind a patient-donor match, and we discuss in more detail the different sources of HSCs and their respective limitations, advantages, costs and benefits. A glossary of acronyms and definitions of medical terms can be found in Section EC.4 of the e-companion.

2.1. Patient-donor histocompatibility

The genetic compatibility (histocompatibility) between a donor and a patient is the basis for a match and for a successful transplant. Histocompatibility is determined by the body's immune system, which uses proteins known as human leukocyte antigens (HLAs) in the cells to distinguish between local and foreign cells. A stem cell transplant is likely to be successful only if the donor's HLA type is sufficiently close to the recipient's, i.e., if the foreign and local cells are similar enough to one another.

To determine whether an HLA-type match exists between a donor and a recipient, until recently the medical standard has been to compare specific gene contents at three locations or loci within chromosome pair six – HLA-A, HLA-B, and HLA-DRB1 – at a low level of resolution.[3] This chromosome pair is made up of two individual chromosomes, one inherited from each parent, with the combination of alleles for HLA-A, HLA-B, and HLA-DRB1 on each individual chromosome known as a haplotype. A patient's phenotype, meanwhile, is determined by combining the gene content inherited from both parents (i.e., taking both haplotypes together). The phenotype is typically represented by the string "$AaBbDd$", where Aa represents the two copies of genetic code (i.e., the two alleles, one from each haplotype) at locus HLA-A, and with Bb and Dd representing the two alleles at HLA-B and HLA-DRB1, respectively. Thus, the histocompatibility match between a BM donor or CBU and a recipient is determined by the degree of similarity between the alleles that make up their phenotypes. If the six alleles mentioned above match between the donor and the recipient, this is referred to as a 6/6 ("six of six") match. If there is a mismatch in one of the alleles, it is called a 5/6 match, while a mismatch in two alleles is referred to as a 4/6 match.

The number of distinct alleles known to exist at each locus is increasing over time as scientists continue testing more individuals in the population. Mori et al. (1997) use a sample of just over 400,000 registered BM donors in the NMDP in 1995 and identify 20 unique alleles at HLA-A, 39 at HLA-B and 14 at BLA-DRB1, meaning that in theory there are at least 10,920 ($= 20 \times 39 \times 14$) different haplotypes and 59.6 million ($= N(N+1)/2$ where $N = 10,920$ – see Sonnenberg et al. 1989) distinct phenotypes that can exist. However, not all possible phenotypes show up in the population. Based on historic phenotyping and subsequent extrapolation, researchers estimate there to be closer to ten million unique phenotypes present

[3] More recently it has been found that outcomes are improved by using higher-resolution matching and by considering at least one additional gene from chromosome pair six. We will discuss the use of more refined matching in Section 7 of this paper.

in the U.S. population (Bergstrom et al. 2009). An estimate of the relative frequencies of these phenotypes and their respective matching probabilities is important for our analysis and is derived later in Section 4.

Note that a critical constraint common to both BM donor recruitment and CBU collection is that the phenotype of a donor or CBU is only discovered after the processing costs are incurred. This means that selective screening based on phenotype is not possible (though selection based on other dimensions, such as race or geography, has been considered in the literature and in practice). While a registry or bank may still decide to reject an incoming BM donor or CBU based on phenotype after screening has occurred, this is rare in practice since the costs associated with adding donor details to the registry and annual storage costs of CBU are low relative to the costs of recruitment, extraction and typing. This inability to perform selective screening therefore makes finding the target annual donor recruitment and CBU collection rates and the corresponding size of the BMR and CBB a primary concern for researchers and policymakers.

2.2. Sourcing and extracting HSCs from bone marrow

Traditionally, the source of HSCs for treating patients with hematological diseases has been BM tissue extracted from the matched donors. (The medical and technical details of the genetics, cell biology and matching process for BM transplants are detailed in Beatty (1994).) As mentioned in Section 1, the search for compatible BM donors is facilitated by BMRs that operate in different countries and maintain information on registered donors. These BMRs are responsible for connecting physicians with potential donors and for recruiting new donors by way of direct advertising, celebrity campaigns, and recruitment drives.

Following recruitment, donor typing (phenotyping) is a reasonably straightforward process. When a consenting potential donor arrives, a simple cheek swab is taken to determine their phenotype. The donor's phenotype is then registered in the BMR together with their contact details. The BMR is therefore simply a database containing the name, demographic information, contact details and phenotype of every registered donor. Whenever a patient requires a stem cell transplantation, a formal search of the BMR is initiated by the transplant center in which they are being treated. If a suitable MUD is found in the registry, then the donor will be contacted to set up a donation. However, on average only approx. 45% of contacts actually result in a donation due to, e.g., difficulties in contacting the potential donor, scheduling conflicts, and ineligibility following medical evaluation (Gragert et al. 2014, Dehn et al. 2019).

If a willing and eligible MUD is found, then stem cells are extracted in one of two ways: 1) a surgical procedure in which liquid bone marrow tissue is extracted from the back of the pelvic bone using a needle under general or local anesthesia, or 2) a non-surgical procedure in which intravenous methods are used to induce the formation of excess stem cells in the donor's bloodstream and those cells are then extracted by way of a blood draw. It is worth bearing in mind that both of these procedures cause great inconvenience to the donor, necessitating, e.g., travel to the medical site, time spent in the hospital, and discomfort during and after the procedure.

2.3. Cord blood as an alternative source of HSCs

In contrast to bone marrow, the advent of cord blood as an alternative source of HSCs is a relatively recent innovation that only started approx. 20 years ago (see Navarrete and Contreras (2009) for a historical overview of CB banking and collection, storage, and release-related factors that are known to impact patient outcomes). CBUs are formed using blood from the umbilical cord and placenta that is collected shortly after childbirth and processed before being frozen in liquid nitrogen. This requires prior consent from the parents, and parents can decide whether to donate the stem cells extracted from the UCB to a public bank, which stores them for public usage, or a private bank, which is a commercial organization where CBUs are stored for potential personal future use. Since not all CBUs are considered usable after extraction, historically approx. 50% of those donated to public banks are discarded for failing to meet minimum standards (Howard et al. 2008).

There are currently 26 public CB banking services in the U.S. (Parent's Guide to Cord Blood Foundation 2020). Each is affiliated with local hospitals, with the maternity wards of these hospitals serving as the main suppliers of CBUs to public banks. This means that CBUs have already been physically collected prior to demand (i.e., prior to a physician searching for a match). As a consequence, unlike BM – which is not collected until there is demand and for which the BMR is essentially a database – CBUs must be physically preserved in cryogenic facilities. Thus, capacity investment and costs associated with collecting, storing and maintaining CBUs for public use are much higher than they are for BM. Their supply is also accordingly limited, with approx. 265,000 units currently listed in the NMDP's Be The Match program (World Marrow Donor Association 2020).

From a matching perspective, CBUs are a highly flexible resource since the phenotype matching requirements are less stringent than they are for BM. In particular, as highlighted in Section 1, while a 6/6 is always preferred, a UCB transplant can be successful even with lower 5/6 and 4/6 matches.[4] (Note that the preference for higher-level matches is due to faster engraftment rates, rather than evidence of differences in survival rates – see Section EC.3 of the e-companion). This flexibility comes from the fact that stem cells from UCB are not fully developed (i.e., the cells are antigen-inexperienced naïve) and hence are more adaptable in the recipient's body. Thus, overall, from an operations perspective, CBUs come at a relatively higher cost than BM donation, but they are also a more flexible resource that is more immediately retrievable.

However, one important limitation of using CBUs is that they cater more to the younger rather than the older population. This is because CBUs typically contain a lower number of HSCs (i.e., a lower total

[4] Note that it is clinically recommended that the two allele mismatches do not occur on the same locus for a 4/6 match. For example, a mismatch on alleles "A" of locus HLA-A and "B" of locus HLA-B is permitted, but mismatches of both alleles "a" and "A" on locus HLA-A is not recommended. In practice, only approx. 5% of UCB transplants have two mismatches at the same locus, so in this paper we follow Gragert et al. (2014) and exclude such matches from consideration.

nucleated cell (TNC) content) than is extracted from BM. As the minimum TNC content requirements increase with patient body weight (and hence indirectly with age – see Section EC.3 of the e-companion for further details), it has been estimated that while most patients under the age of 20 will be able to find at least a 4/6 match in public CBBs that exist in the U.S. today, for adults this drops to around 80% (Gragert et al. 2014). In comparison, an equivalent BM transplantation will provide recipients with approx. ten times the TNC content of UCB. Partially because of this difference in TNC content, patients receiving CBUs are also at increased risk of graft failure and slower rates of engraftment (complications that are described in Section EC.3 of the e-companion). As a result, in most cases, physicians prefer BM over UCB for allogeneic transplants. However, there is no guarantee that a perfect 6/6 match will always be found in the BMR. Hence, despite the limitations of CBU transplantation, its greater flexibility can potentially play a major role in catering to the needs of patients with rare phenotypes who cannot find a match in the BMR.

With a deeper understanding of the biological and operational features of both sources of HSCs, we next proceed to review the literature relevant to the problem at hand.

3. Literature Review

Our paper primarily builds on three streams of literature: 1) medical papers that discuss the biology associated with finding a suitable match, 2) economic analyses that focus on improving social welfare by allocating resources for transplants more efficiently, and 3) operations management (OM) literature relating to capacity decisions in contexts where alternative resources exist which are partial substitutes.

3.1. Medical literature and economic analyses

An extensive body of medical literature discusses the biology behind finding a match, evaluates the relative advantages and disadvantages of BM and UCB as sources of HSCs, provides recommendations as to which source to choose and when, and compares relevant patient outcomes following allogeneic transplants depending on source and degree of matching (see, e.g., Giralt and Bishop 2009, Kekre and Antin 2014, Dehn et al. 2019). Much of this is important in modeling the joint optimization problem and is discussed in detail in Sections 2, 4 and 5.

Also relevant to this paper is medical literature that seeks to derive the probability of finding a patient's histocompatible match. Sonnenberg et al. (1989) were the first to formally calculate the probability of finding a MUD for a given BMR size. Using updated data, Mori et al. (1997) study the haplotype frequencies of the North American population and provide guidance for matching and suggest donor recruitment strategies. These works are extended by Beatty et al. (2000), who mathematically calculate the probability of a 5/6 match, and by Kollman et al. (2004), who assess how increasing donor recruitment (especially targeted recruitment by race) affects matching probabilities. More recently, Gragert et al. (2014) estimated that with

the existing BMR and public CBB supply, the likelihood of a patient finding a suitable match ranges from 91 to 99% depending on patient age and ethnicity.

Building on the medical literature and using the biological underpinnings as inputs, a separate stream of literature analyses the problem of finding the optimal BMR or CBB size from a cost-benefit perspective to inform policy. Bergstrom et al. (2009), for example, model and optimize the size of the BMR and estimate the altruistic benefit associated with being a BM donor depending on donor race. Meanwhile, Howard et al. (2008) build a mathematical model to show how, given a fixed demand for CBUs, cost-effectiveness reduces as the size of the CBB increases. Meanwhile, in an unpublished manuscript, Fève and Florens (2010) develop separate mathematical models for the BMR and CBB to find the matching probabilities associated with each institution when used in isolation. Specifically, after simulating over the phenotype distribution in France, the authors assess how the matching probabilities change when using either a BMR or a CBB of varying size (with size incremented in fixed intervals). However, the analysis is primarily descriptive; it does not optimize over the respective sizes of the two institutions and does not optimize jointly.

A major limitation common to these existing studies is that when the size of the BMR or CBB is varied, the presence of the other source is ignored. Since this approach will clearly lead to sub-optimal policy recommendations, in this paper we instead consider the problem of jointly optimizing the size of the BMR and CBB, recognizing that these two sources of HSCs are not perfect substitutes (e.g., they differ in their matching criteria, costs, and medical effectiveness). However, considering the two institutions jointly introduces a technical challenge: As the size of one institution changes, it changes the distribution of demand that spills over to the other; consequently, the distribution of supply also changes. Over time, the distribution of demand and supply therefore become increasingly misaligned, resulting in fewer matches. The existing literature does not capture these dynamics. In contrast, our study finds these dynamics to have important policy implications, as we show by explicitly accounting for temporal-variation in the matching probabilities in a simulation study and comparing it to a static model (under restrictive conditions).

3.2. Operations management

As noted earlier, CBUs can be viewed as both a flexible resource and as an (imperfect) substitute for BM. Our paper therefore touches on three important streams of literature drawn from the OM domain: 1) capacity decisions with flexible resources, 2) inventory management with substitute products, and 3) matching and resource allocation in the kidney exchange market. We discuss our work in relation to these streams below.

The recent review article by Song et al. (2020) provides a thorough survey of OM literature relating to capacity and inventory management in general, and so our focus here is on discussing those papers most related to our work. In terms of capacity planning, Van Mieghem (1998) provides guidance for optimal

investment in dedicated and/or flexible capacity given two different products that differ in aspects such as cost, price and demand uncertainty. This is extended to multiple products and multiple resources in the work of Harrison and Van Mieghem (1999). A more detailed literature review on the early work relating to strategic capacity management and hedging under demand uncertainty can be found in Van Mieghem (2003). In the body of literature concerning inventory management with substitute products, Bassok et al. (1999) focuses on single period multi-products with substitutability, proving that a greedy policy of inventory replenishment is optimal in this case. Lee et al. (2016) demonstrate the value of utility-based choice models in planning inventory for new and used textbooks with stockout-based substitution. Meanwhile, Schlapp and Fleischmann (2018) derive optimal inventory policy for a capacity-constrained firm selling multiple partially substitutable products over a finite season.

Also closely related to our paper is literature discussing the problem of matching and resource allocation in the kidney exchange market. Prior to the seminal work of Roth et al. (2004, 2005), Zenios (2002) modeled the kidney exchange program as a double-sided queue matching problem and sought to find policies to maximize the expected total discounted quality-adjusted life years (QALYs). Building on this work, Su and Zenios (2006) studied the efficiency-equity trade-off in kidney transplantation and designed a mechanism to elicit truthful private information from patients. Glorie et al. (2014), meanwhile, focused on the market clearance problem and on finding a donor-to-patient allocation that is optimal with respect to multiple criteria. Many other contributions of the OM community to improving the kidney exchange process are summarized in Anderson et al. (2015). Progress on this problem has continued, with recent work, e.g., Ding et al. (2018), Dickerson et al. (2019), Blum et al. (2020), taking an algorithmic approach to improving the efficiency of the kidney exchange market.

To the best of our awareness, Gökalp et al. (2020) is the first paper in the operations management literature related to the broad theme of HSC transplants. The paper focuses on reducing delays in transplantation after matching by building advanced testing capacity and coordinating rapid donor mobility. Their work highlights some of the structural issues that prevent the actual transplant rate from reaching the theoretical match rate suggested in the literature. We complement this work by establishing the first-best solution from a system design perspective, which acts for a benchmark for literature focusing on operational efficiencies.

While at a high level there are similarities between our problem and those discussed in these related streams of OM literature, our problem has several features that have not previously been studied. First, as the capacity of one resource (e.g., the BMR) changes, it affects not only the total demand spilling over to the other resource (e.g., the CBB) but also the distribution of that demand. This introduces an additional inter-dependency that must be accounted for when making capacity decisions. Second, as previously noted, the composition of inventory changes dynamically over time due to the spillover effect. Therefore, a static

single-period model is insufficient to capture this dynamic, and a simulation-based approach is required to identify the optimal allocation decision. Third, the inventory management problem described in our paper is unique in the sense that we have more than 10 million different possible SKUs for BM and CB. The small inventory size relative to the wide variety of SKUs implies that the supply (i.e., the composition of the SKUs held in inventory) at any point in time is a random sample. Moreover, the composition of this random sample changes dynamically with time and cannot be deterministically controlled (since replenishment also occurs randomly), making sample composition highly dependent on the capacity decisions made by the policymaker. These factors distinguish the problem analyzed in this paper from those studied in the extant OM literature (see Song et al. 2020).

4. Data Preparation

The data used in our study contains the population frequency distribution of phenotypes compiled by Mori et al. (1997), which is based on a sample of about 400,000 individuals who were registered with the NMDP in 1995 and whose HLA-A, -B, -DR phenotypes were recorded. To obtain the probability of every phenotype in the population, the phenotypes (i.e., "$AaBbDd$") are split into their constituent haplotypes (i.e., "ABD" and "abd"). Mori et al. (1997) then calculated the frequency of every haplotype h_{ABD} using the Expectation-Maximization (EM) algorithm prescribed by Dempster et al. (1977). Based on these haplotype frequencies, we construct an exhaustive list of all phenotypes and their frequencies in the population according to the laws of genetics. Specifically, according to the Hardy-Weinberg equilibrium, the population frequency for phenotype $AaBbDd$ is given by[5]

$$p_{AaBbDd} = 2[h_{ABD}h_{abd} + h_{ABd}h_{abD} + h_{AbD}h_{aBd} + h_{aBD}h_{Abd}] \tag{1}$$

This equation reflects the fact that any of the individual terms on the right side of Equation (1) can give rise to the same phenotype, and the multiplicative factor 2 is to account for the equiprobable mirror image haplotype combinations, which can be inherited in two ways, i.e., $h_{ABD} * h_{abd} = h_{abd} * h_{ABD}$. We iterate the above equation for all possible phenotype combinations and obtain 10,804,408 possible distinct phenotypes for the U.S. population. Thus, p is a vector of size 10,804,408. The probabilities of the phenotypes vary from 10^{-3} to 10^{-13}.

Next, we calculate the probabilities of "6 of 6" (6/6), "5 of 6" (5/6), and "4 of 6" (4/6) matches for each of the phenotypes in the vector p. These are required to calculate the likelihood of a patient with a particular phenotype finding a match in the BMR or CBB. Let m_k^{AaBbDd} be defined as the probability that

[5] Note that the Hardy-Weinberg equilibrium maintains the assumptions that individuals mate with others of their own race and that mating within each race is random with respect to HLA type.

a randomly chosen member of the population is at least a $k/6$ match for phenotype $AaBbDd$. Clearly, $m_6^{AaBbDd} = p_{AaBbDd}$. For a 5/6 match we find

$$m_5^{AaBbDd} = p_{.aBbDd} + p_{A.BbDd} + p_{Aa.bDd} + p_{AaB.Dd} + p_{AaBb.d} + p_{AaBbD.} - 5p_{AaBbDd}, \tag{2}$$

where the "." symbol in any phenotype subscript denotes summation over all possible values in the corresponding position. It is easy to see that m_5^{AaBbDd} is nothing but the sum of probabilities of phenotypes that have a mismatch with $AaBbDd$ in only one of the six places, with a final subtraction term that accounts for double-counting of $AaBbDd$.

The mathematical expression for calculating a m_4^{AaBbDd} is not as straightforward and, to the best of our knowledge, has not previously been derived in the literature. The additional complexity is due to the added constraint that for a 4/6 match to succeed, there must be at least one matching antigen for each of the three pairs Aa, Bb and Dd (see Footnote 4). Combinatorics gives us $\binom{6}{2} - 3 = 12$ such possible sequences.

Proposition 1 *The probability of at least a* $4/6$ *match for phenotype* $AaBbDd$ *is given by*

$$\begin{aligned} m_4^{AaBbDd} = {} & p_{.a.bDd} + p_{.aB.Dd} + p_{.aBb.d} + p_{.aBbD.} + p_{A..bDd} + p_{A.B.Dd} \\ & + p_{A.Bb.d} + p_{A.BbD.} + p_{Aa.b.d} + p_{Aa.bD.} + p_{AaB..d} + p_{AaB.D.} \\ & - 3m_5^{AaBbDd} - 8p_{AaBbDd} \end{aligned} \tag{3}$$

where again the "." symbol in any phenotype subscript denotes summation over all possible values in the corresponding position.

The detailed derivation of Equation (3) is provided in Section EC.1 of the e-companion.

In Figure 1 we plot histograms for the m_6^{AaBbDd} (bottom), m_5^{AaBbDd} (middle) and m_4^{AaBbDd} (top) matching probabilities. Specifically, the x-axis gives the probability that a randomly chosen member of the population will be at least a $6/6$ (bottom), $5/6$ (middle) or $4/6$ (top) match for phenotype $AaBbDd$, while the height of each bin specifies the number of phenotypes (out of the 10,804,408 possible) that fall into each probability range. From Figure 1, it is evident that the median probability of matching with a randomly chosen member of the population improves as the matching criteria is relaxed (from 6/6 to 4/6). Furthermore, the dispersion of probabilities narrows and concentrates around the median. This figure thus demonstrates the value that CBUs can bring to the system by enabling lower-level (i.e., 5/6 and 4/6) matches.

5. Mathematical Model

In this section, we first develop a model to predict the performance of the BMR and then use this model to address the question of the optimal registry size. Next, we consider the performance of a system consisting of both a BMR and a CBB and discuss the challenges of mathematically modeling this joint system.

Figure 1 **Distribution of m_4^{AaBbDd} (top), m_5^{AaBbDd} (middle), and m_6^{AaBbDd} (bottom) matching probabilities across phenotypes.**

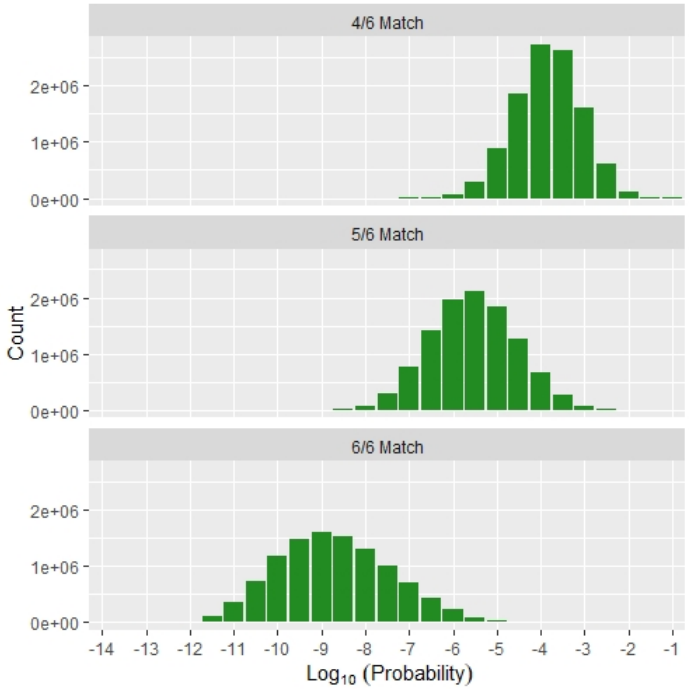

5.1. Bone marrow registry model

As discussed in Section 2.3, donated BM is generally preferred to CBUs as a way to satisfy the demand for allogeneic transplants due to the lower odds of graft rejection and higher TNC content as compared to CBUs. Moreover, maintaining a public CBB requires cryogenic storage facilities and is expensive. Consequently, many emerging economies do not have CBBs and instead rely on BMRs to find healthy donors. We therefore start by assessing the performance of a stand-alone BMR. The model we develop here will later also be used to analyze the joint BMR/CBB system.

We start by setting B as the number of registrants in the BMR, which is assumed to be held constant. We also assume that the composition of the registry always reflects the population phenotype distribution p, i.e., that recruitment of potential donors and attrition are independent of the registrant's phenotype.[6] This means that at any time, the probability that a randomly chosen registrant has phenotype i is equal to p_i (as defined in Equation (1)). Meanwhile, the probability that a random patient will arrive and place a request for phenotype i is also given by p_i.

Although there are B registrants in the BMR, this number is effectively reduced by the fact that registrants may not be willing or able to donate when called upon to do so, for reasons discussed in Section 2.2. Therefore, we assume that a registrant will "show up" when requested with probability δ, independent of all

[6] It is simple to relax this assumption and allow the phenotype distribution in the BMR to follow some other stable distribution.

other requests made to this or any other registrant and independent of their phenotype. (Note that following a donation request, this registrant will remain on the registry regardless of whether they show up or not.)

Now, let $\overline{M}$ be the event that an arbitrary request for a donation cannot be met from the registry, and let T be the phenotype associated with the request. Then

$$\Pr\{\overline{M}|T=i\} = (1-\delta p_i)^B, \tag{4}$$

reflecting the fact that for every registrant in the registry, the joint probability of having type i and showing up equals δp_i. This means that

$$\varphi(B) = \Pr\{\overline{M}\} = \sum_i \Pr\{T=i\} \Pr\{\overline{M}|T=i\} = \sum_i p_i(1-\delta p_i)^B \tag{5}$$

is the fraction of donation requests that remain unsatisfied in a BMR of size B.

The primary objective of the model is to maximize the net benefits derived in the process of matching patients with donors in the BMR.[7] It must therefore account for the total value of lives saved, the cost incurred in establishing the registry, and the cost of replenishing the registry to compensate for attrition. The key parameters for this model and estimates of their values are:

- The annual demand for unrelated HSC transplants, which is reported to be close to $\lambda =$13,000 patients per year in the U.S. (Besse et al. 2015, Health Resources and Services Administration 2020b).
- The probability distribution of the different phenotypes in the population, p, which was estimated in Section 4.
- The time horizon over which to optimize, which we set equal to ten years. In particular, stem cells are currently being tested for application in treating a wide variety of diseases and disorders (Aly 2020) and the technology associated with extracting, storing, and administering them is rapidly evolving (Brown et al. 2019). These, and other as yet unforeseeable changes, may drastically affect future demand for HSCs and, consequently, the recommended size of the BMR/CBB. Our choice to present the results for a 10-year horizon is thus a conservative one (see also Section 6.4 of the paper where sensitivity analyses are performed using different levels of demand).

[7] As mentioned earlier, our analysis aims to achieve the first-best solution from a system design perspective, hence we follow the convention in the literature and assume that an eligible and willing match always leads to a transplant (e.g., Bergstrom et al. 2009). However, aside from no-shows, which we account for, there may be various other reasons why a transplant does not occur despite there existing an available match in the BMR or CBB. For example, the patient may die before the donor or CBU arrives, the patient may be unable to afford the cost of treatment, or the patient may not be willing or able to wait, and so may opt for a haploidentically mismatched related donor (i.e., a 3/6 match) instead. Consequently, the number of allogeneic transplants performed in the U.S. annually (at least 5,000 (Health Resources and Services Administration 2020d)) will not tally with the number expected when using the theoretical match rate. Our analysis therefore serves as a benchmark for future research focusing on addressing these operational inefficiencies.

- The annual attrition rate, α. This is estimated by comparing the average age at which a potential donor joins the BMR – given by Bergstrom et al. (2009) as 35 – with the age at which they "age out" of the registry – taken to be 61, the age at which patients are officially removed from the Be The Match registry because donation poses increasing risks to their health. Thus, on account of age censoring, the annual rate of attrition is set at $\alpha = 0.04$ (donors spend on average 25 years on the registry).
- The no-show rate, δ, which was discussed in Section 2.2 and reported to be approx. $\delta = 0.45$.
- The value of a match through the BMR, v_b. We follow Bergstrom et al. (2009) and assume a value of \$1.2 million. This estimate is based on the concept of the "value of a statistical life" (VSL), which is the marginal rate of substitution between survival probability and wealth. Bergstrom et al. (2009) use a VSL of \$6.5 million and estimate a 21% increase in the likelihood of long-term survival following a successful transplant (relative to a patient who does not receive a transplant). After subtracting the median hospital cost of the transplant, \$166,000, the value of a match is thus estimated as $v_b = (\$6.5 \times 0.21 - \$0.166) \times 10^6 \approx \$1.2 \times 10^6$.
- The cost associated with recruiting a donor, c_b, which, following Bergstrom et al. (2009), is equivalent to the cost of phenotyping and equal to $c_b = \$105$ per registrant.
- The annual costs of replenishing the registry to compensate for attrition so as to maintain it at a constant size of B, which is equal to $\alpha B c_b$.
- The cost of establishing the BMR, which we assume to be equal to the total typing costs Bc_b. This cost is amortized over ten years, i.e., the time horizon chosen for this study, as noted above.[8]

Incorporating these parameters, our objective is to solve the problem $\max\{f(B)|B \geq 0\}$ where

$$f(B) = \lambda(1 - \varphi(B))v_b - (\alpha + 0.1)Bc_b. \tag{6}$$

Equivalently, we solve the problem $\min\{g(B)|B \geq 0\}$ where

$$g(B) = \lambda v_b - f(B) = \lambda\varphi(B)v_b + (\alpha + 0.1)Bc_b. \tag{7}$$

In Section EC.1 of the e-companion, we show

Proposition 2 *g(B) is convex in B.*

Note that our formulation does not include discounting. This is because, first, discounting the value of a life saved is questionable. Second, it is not unreasonable to assume that inflation and discounting for the out-of-pocket costs cancel each other out. Third, discounting merely complicates the presentation and does not change the insights and conclusions. However, in Section EC.5 of the e-companion we show the impact of incorporating discounting into the joint optimization model.

[8] We could, alternatively, only consider the typing costs associated with increasing the size of the registry to B from its current level (8 million potential donors in the U.S. – see Footnote 2). Doing so produces identical results.

5.2. Numerical investigation of the BM registry model

First, note that an analytical solution is not available. Fortunately, the function $g(B)$ is convex in B, so numerical optimization is straightforward, even though the size of the vector p is vast (10,804,408$\times$1) and the decision variable is in the exponent of each of the terms. In Figure 2a we plot the region around the optimal value of B, which clearly depicts the convex structure of the function $g(B)$.

We find that the optimal value of the registry size B^* for an annual attrition rate of $\alpha = 0.04$ and show-up rate of $\delta = 0.45$ is 34×10^6, meaning that the registry should have ± 34 million potential donors. With the optimal registry size suggested by our model, on average 12,405 out of the 13,000 patients in a year find a match, leading to an expected fill rate of 95.42% and saving 2,605 lives in expectation (12,405$\times$0.21). Donor recruitment must be sufficient to maintain the registry at 34 million donors. With a 4% attrition rate, this means that approx 1.36 million new donors must be recruited annually in the U.S. alone. This is close to the total number of donors added every year to the Be The Match registry globally between 2017 and 2019 (Health Resources and Services Administration 2020a). Moreover, the registry is currently estimated to have only 11.5 million donors i.e., only $\approx$ 34% of the recommended level. This gap highlights the magnitude of the efforts needed to expand the registry size and maintain it at 34 million donors. Further, the annual cost to recruit new donors to compensate for attrition is estimated at $143 million, while the amortized fixed cost for expanding the registry beyond its 11.5 million current registrants is $238 million. Thus, in this model in which the BMR is the sole source to address the needs of 13,000 patients per year, the total annual cost amounts to $381 million.

Our recommendation of 34 million BM registrants is approx. 13 million more than the 21 million prescribed by Bergstrom et al. (2009). This is primarily due to a difference in how we account for attrition as well as the addition of no-show rates in our model. Specifically on the first point, Bergstrom et al. (2009) arbitrarily assume that the match rate (i.e., % of patients who find a match in the BMR) declines by 2% per year after the first year. This can be explained by the fact that recurring donor recruitment does not occur in their model. Consequently, over time, the size of their BMR effectively reduces due to donor attrition, and hence the match rate also decreases. However, a 2% decline in matches per year leads to a large number of additional lives lost, dampening the benefits of the BMR in Bergstrom et al. (2009) and resulting in a significant underestimate the optimal BMR size. By comparison, we assume (more realistically and consistent with practice) that if there is donor loss then new donors will be recruited to replace them (at a cost), and therefore that the match rate instead remains stable over time.

Note that the values of the attrition rate (α) and show-up rate (δ) may vary over time (e.g., if younger and healthier donors are recruited). Therefore, it is helpful to understand how sensitive the optimal BMR size, B^*, is to changes in these parameters. First, we investigate the impact of the attrition rate α for a given

Figure 2 (a) Structure of $g(B)$.
Sensitivity of the optimal registry size (B^*) against (b) attrition rate α and (c) show-up rate δ

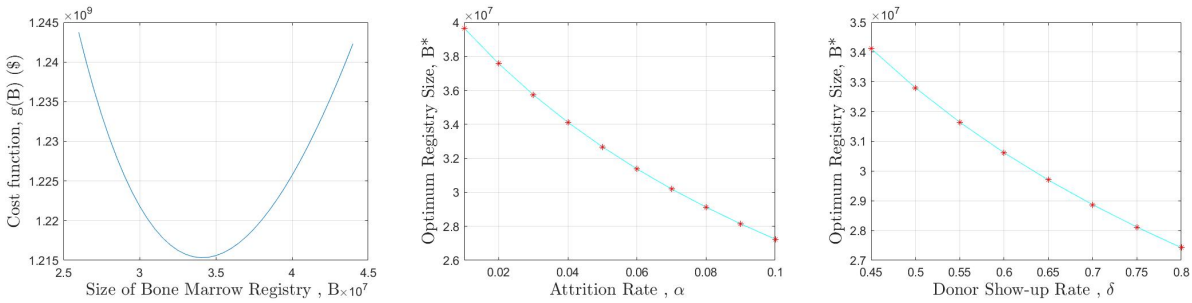

(a) (b) (c)

value of the show-up rate (45%). Figure 2b shows the steep decrease in the optimal BMR size, B^*, as the annual attrition rate α is increased. The decrease is due to the higher annual recruitment cost incurred with higher attrition rates. In turn, this also impacts the number of matches, which decreases when there are fewer BM donors in the registry. Figure 2c, meanwhile, shows that the optimal BMR size (B^*) decreases as the show-up rate δ increases. The decreasing optimal size is a consequence of more registrants showing up to donate when called upon, which means that fewer registrants are required in the BMR. In this case, although the registry size decreases, the number of lives saved remains almost constant (as the increased show-up rate compensates for the decrease in registry size).

As mentioned at the beginning of Section 5.1, our model provides a robust framework for use in emerging economies where public CBBs are tough to establish. After the population's phenotype distribution is estimated, the framework above can be utilized to prescribe an optimal BMR size in these locations. Next, we turn our attention to CBUs and explore how much value they bring to the system.

5.3. A joint model

CBUs as an alternate source of HSCs have significant potential value for patients requiring allogeneic transplants, primarily due to greater flexibility in their matching requirements (see Section 2.3). However, despite this advantage, the probability of graft rejection is higher for UCB than for BM transplants, as mentioned earlier. Therefore, BM donation is almost always the doctor-preferred option, and when a transplant request comes in, the standard procedure usually involves searching for a BM donor first. If none are available, a match is sought in the CBB in the preference order of 6/6, 5/6 and 4/6. When a match is found in the CBB, the TNC level in the CBU must be evaluated to assess whether it is sufficient for transplantation. As noted in Section 2.3, this depends primarily on a patient's weight (and so it also depends, indirectly, on their age). A request for a CBU is thus associated with a minimum TNC count requirement, so not every matching unit can be used to fulfill a given request.

It is essential to understand what happens to the compositions of the BMR and CBB if this search process is repeated over time. In particular, to obtain a computationally tractable joint optimization problem, we need a tractable model to predict CBB performance. This turns out to be problematic for several reasons. First, the composition of even a "new" CBB, freshly sampled from the population distribution, does not reflect the demand placed on it, because the CBB only caters to the overflow from the BMR. For a CBU with a given phenotype p_i, there can be demand from patients with several other phenotypes (i.e., 5/6 and 4/6 matches). The TNC content requirement complicates things further: As we will shortly show, there are $16 \times 10{,}804{,}408$ distinct possible combinations of phenotype-TNC content.

Second, the composition of the CBB changes over time: CBUs with an "easy to match" phenotype will be consumed faster and will tend to disappear from the CBB, while units with "hard to match" phenotypes (stale inventory) will accumulate. A similar effect will tend to cause low TNC content units to remain in the CBB, while high TNC content units will be consumed and disappear. Hence if the CBB's size is held constant, its performance will deteriorate over time. However, since in our model part of the CBB is replaced each year with new units from the population (in reality this is a continuous process), this will cause the deterioration in performance to stabilize after a number of years.

Hence, a direct analytical approach appears to be a dead end. We therefore take two different routes around this difficulty. To begin, we analyze an "extreme policy" where we assume that the (probabilistic) composition of the CBB always reflects the population distribution p and does not change over time. This policy could be approximated in practice by replacing the entire CBB each year.[9] Then, in Section 6, we use a simulation study of CBB performance over time to develop a predictive regression equation. This regression model is then combined with the BM performance model to predict the performance of a joint BMR/CBB system, and we use this to find the optimal registry and bank sizes for such a joint system.

5.4. Extreme policy

To analyze the performance of the CBB under the extreme policy, we first note that the expected annual demand for CBUs is $\lambda\varphi(B)$ where, following Section 5.1, λ is the annual demand for unrelated HSC transplants and $\varphi(B)$ as defined in Equation (5) is the probability that an arbitrary patient's demand cannot be met by a BMR of size B. Let q_i be the probability that a demand for a CBU is for phenotype i. Note that generally $q_i \neq p_i$, since it is only phenotypes that are not present in the BMR that will be demanded from the CBB. As in Section 5.1, we again define T as the phenotype of an arbitrary patient and $\overline{M}$ to be the event that an arbitrary patient's demand is not met by the BMR. Hence

$$q_i = \Pr\{T = i | \overline{M}\} = \frac{\Pr\{T = i \wedge \overline{M}\}}{\Pr\{\overline{M}\}} = \frac{p_i(1-\delta p_i)^B}{\varphi(B)}. \tag{8}$$

[9] The choice of one year is arbitrary and any other time interval could have been chosen. The goal of analysing the extreme policy, however, is simply to demonstrate the additional value of a CBB and to introduce features that will help us in Section 6 to model a system that more closely resembles reality. This is not a policy that we would recommend or expect to be implemented in practice.

Turning to the composition of the CBB, let C be the number of CBUs in the bank and η be the TNC content of an arbitrary CBU. Gragert et al. (2014) estimates the distribution of adequate TNC content proportions available in public banks for adults and children based on weight deciles. Using this, we model the TNC content requirements by distinguishing 16 different levels of TNC content for CBUs (i.e., $\eta \in \{1, \cdots, 16\}$). We then let $\zeta \in \{1, \cdots, 16\}$ specify the minimum TNC content required by a random request for a CBU and assume that η and ζ are independent of each other and of the phenotype needed. The probability distributions of η and ζ are given in Table 1.

TNC level (k)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Sum
$\Pr\{\zeta = k\}$	.06	.02	.02	.02	.08	.02	.02	.08	.08	.10	.08	.08	.08	.10	.08	.08	1.00
$\Pr\{\eta = k\}$	.17	.17	.29	.09	.10	.04	.01	.04	.02	.02	.01	.01	.01	.01	0.005	0.005	1.00

Table 1: Probability distributions of TNC content requested (ζ) and available in CBUs (η).

If the CBB composition reflects the population distribution, the probability that an arbitrary CBU can satisfy a demand for phenotype i is given by $\Pr\{\eta \geq \zeta\} m_i^4$, where m_i^4 is the probablity of at least a 4/6 match for phenotype i (see Equation (3)), and $\Pr\{\eta \geq \zeta\}$ is the probability that a CBU has sufficient TNC content to satisfy the demand. Hence the probability that a demand for phenotype i and a TNC content of at least k cannot be satisfied by any unit in the CBB is given by $(1 - \Pr\{\eta \geq k\} m_i^4)^C$. So the probability that an arbitrary demand for phenotype i cannot be satisfied by any unit in the CBB equals

$$\psi_i = \sum_{k=1}^{16} \Pr\{\zeta = k\} (1 - \Pr\{\eta \geq k\} m_i^4)^C. \tag{9}$$

This means that the probability that an arbitrary demand on the CBB cannot be satisfied equals $\sum_i q_i \psi_i$, and so the probability that a random patient cannot find a match in either the BMR or the CBB thus equals

$$\psi(B, C) = \varphi(B) \sum_i q_i \psi_i = \sum_i p_i (1 - \delta p_i)^B \left[\sum_{k=1}^{16} \Pr\{\zeta = k\} (1 - \Pr\{\eta \geq k\} m_i^4)^C \right]. \tag{10}$$

For the joint BMR and CBB system, the expected annual benefit of matches made is therefore equal to $\lambda(1 - \varphi(B)) v_b + \lambda(\varphi(B) - \psi(B, C)) v_c$. Note that v_c is calculated in a similar way as v_b, except that due to graft failure the post-transplant long-term survival probability decreases from 21% to 16%. Additionally, the median hospitalization cost for a UCB transplant is estimated to be lower at \$137,000 (Majhail et al. 2009), while the value of statistical life saved remains \$6.5 million. Therefore, the value of a UCB transplant is estimated as $v_c = (\$6.5 \times 0.16 - \$0.137) \times 10^6 = \$903,000$.

In terms of costs, in addition to the annual operating cost of the BMR, which remains equal to $(\alpha + 0.1) c_b B$, we must also account for the costs of operating the CBB. In particular, two additional costs must

be included: the cost of adding a unit of CB to the bank, c_c, and annual holding costs per unit, h. After accounting for the fact that 50% of CBU collected are deemed unfit during processing and are discarded, the effective cost of adding a single CBU to the CBB is given by $c_c = \$2000$, while the annual storage cost is $\$50$ (Howard et al. 2008). Thus, the annual operating cost of the CBB equals $(c_c + h)C$, reflecting the fact that the entire CBB is renewed each year and storage costs are incurred over the entire bank.

Hence, our objective is to maximize the function

$$F(B,C) = -0.1Bc_b + \begin{bmatrix} \lambda(1-\varphi(B))v_b + \lambda(\varphi(B) - \psi(B,C))v_c \\ -\alpha B c_b - C c_c - Ch \end{bmatrix} \tag{11}$$

This is of course equivalent to minimizing the function

$$G(B,C) = \lambda v_b - F(B,C) = \begin{bmatrix} \lambda\varphi(B))(v_b - v_c) + \lambda\psi(B,C)v_c \\ +\alpha B c_b + C c_c + Ch \end{bmatrix} + 0.1Bc_b \tag{12}$$

In Section EC.1 of the e-companion, we show

Proposition 3 *$G(B,C)$ is jointly convex in B and C.*

5.5. Analysis of the extreme policy

Since a closed-form solution cannot be obtained, we numerically minimize $G(B,C)$. The global minimum of the cost function is attained at $B^* \approx$ 19,680,000 and $C^* \approx$ 66,300. Given demand of 13,000 patients, this policy, at its optimum, fulfills the needs of 12,134 patients via the BMR and 579 via the CBB. In other words, the BMR serves 271 fewer patients than it does under the BM only system, but the extreme policy serves 308 more patients overall, increasing the fill rate from 95.4% to 97.8%. Recall, however, that the value of matching via the CBB is lower than the value of matching via the BMR due to the lower survival rates following UCB transplantation. Regardless, this extreme policy still saves 35 more expected lives than with BMR only.

Not only are more lives saved under the extreme policy, but it is also worth noting that the optimal size of the BMR has almost halved in the presence of the CBB. As a consequence, annual operating costs are also lower. The estimated costs of running the BMR and CBB per year are $83 million and $136 million, respectively. The 10-year amortized cost of expanding the BMR to reach the optimum (not required for the CBB, which is renewed entirely each year) adds a further $86 million per year. This results in a total annual operating cost of $305 million, i.e., 80% of the cost of operating the BMR only system.

6. Simulation Study and a Predictive Model to Analyze the Temporal Variation

The results of the mathematical models developed in the previous section provide a strong case for the use of CBUs as a secondary source of HSCs. However, the extreme assumption that all unused CBUs will

be discarded at the end of each year is made purely for tractability reasons. In reality, CBUs often have a shelf-life of approx. 20 years. This introduces a modeling challenge. Even when a CBB initially reflects the population distributions of phenotypes and TNC content, the ongoing consumption of CBUs and their random replacement from the donor population means that the CBB composition will gradually change and become less effective in supplying matches over time. This would be the case even if the demand for CBUs were according to the population distribution, only 6/6 matches were allowed, and the TNC content were not a factor. The effect is exacerbated further, though, by the common occurrence of 5/6 and 4/6 matches, the need for sufficient TNC content in CBUs, and the fact that the demand on the CBB is for those phenotypes that could not be matched by the BMR. Since we are not able to calculate the performance of the CBB with these effects in play, we instead introduce a simulation-based modeling approach.

6.1. Simulation routine

Since an analytical model is available for the BMR, the simulation model focuses on the CBB.[10] The main objective of the simulation is to predict the number of matches that the CBB can make over ten years given a constant CBB size.

At the start of the simulation, the CBB is populated with C independent CBUs drawn from the population distribution p, supplemented with random TNC levels distributed as η. Every year, the CBB receives D requests that could not be filled by the BMR. We assume that for a given BMR size B and no-show rate δ, the probability distribution of the phenotypes of these D requests is given by q (see Equation (8)), while the minimum TNC levels required are distributed as ζ. Since the CBUs are cryogenically preserved in the public banks for a period of 20 years (Howard et al. 2008), they therefore expire at a rate of 5% per year. Hence at the end of each simulated year, $0.05C$ randomly chosen units are removed from the CBB, and sufficient randomly generated CBUs are added to restore the bank to size C. After ten simulated years, the simulation ends and reports the total number of matches, N, and the quantity of unmatched demand. A pseudo-code version of the simulation algorithm is given in Algorithm 1, in Section EC.2 of the e-companion.

The main parameters that are varied in the simulation study are B, δ, C and D. Further inputs are the vectors p, q (which is a function of B, δ and p) and the distributions of η and ζ given in Table 1. The parameter values used in the simulation study were as follows: $B \in \{8, 10, \ldots, 20\} \times 10^6$, $C \in \{0.5, 1, 1.5, 2, 2.5, 3, 4, 5, 6, 7\} \times 10^5$, $D \in \{500, 1000, \ldots, 3000\}$, and $\delta \in \{0.25, 0.35, \ldots, 0.75\}$. All combinations were simulated once, yielding a total of 2,520 cases.

It is evident from the simulation output that, as expected, the number of matches decreases over time. On average over a 10-year period, the number of CBB matches is 11.2% less than what would be expected

[10] This considerably streamlines the simulation; an alternative would be to simulate the composition of the BMR in detail.

from a CBB of the same size that always reflects the population distribution in phenotype and TNC content, with a minimum loss of 2.2% and a maximum of 32.2%. A large CBB size C and a low value of demand D result in the smallest performance degradation, a small CBB and high demand in the largest.

6.2. Predictive model

We now use the output from the simulation to develop a predictive regression model to capture the effect of the decline in the number of matches. We start by defining the key variables in the regression.

First, we note that by using Equations (5) and (10), we can calculate $\tau = 1 - \psi(B, C)/\varphi(B)$ from known quantities. As defined, τ is the fraction of demand on the CBB that can be met by the CBB when it is fresh (i.e., when it reflects the population distribution). In other words, τ is approx. equal to the expected proportion of demand that finds a match in the CBB in the first year of the simulation, with the corresponding expected number of CBB matches τD. The parameter τ is therefore helpful in anchoring the regression to a starting value and in predicting the number of matches over the 10-year period. Furthermore, $10\tau D$ approximates the total matches over a 10-year period when using a CBB that is refreshed entirely each year. Thus, $10\tau D$ can be thought of as an upper bound on the total number of matches in a CBB that loses efficiency over time (i.e., N), while $Y = N/(10\tau D)$ captures the average CBB efficiency over ten years, where starting efficiency (with a fresh CBB) is 100%. In particular, Y can be interpreted as the simulated number of matches into the expected number of matches that would be achieved if an "always fresh" CBB were available. This means that the more the composition of CBB falls out of alignment with the demand placed on it, the lower the value of Y will become. Figure 3a visually portrays the annual maximum, average and minimum fraction of demand matched from the CBB in the 2,520 simulated cases.

In selecting the main variables in the regression, we experimented extensively with various transformations of the simulation inputs (i.e., B, δ, C and D) as regressors. Moreover, we tried several transformations of N, the total simulated number of CBB matches, for the dependent variable. We finally decided to use $Y = N/(10\tau D)$ as the dependent variable and τ, $\ln(C)$, $\ln(\delta B)$, D and D/C as explanatory variables. The final model had an R-squared value of 0.98905 (adj. R-squared = 0.98903), all coefficients were extremely significant (all p-values $< 10^{-50}$), with the predictive equation

$$\hat{Y} = 0.7060 - 0.2443\tau + 0.0657\ln(C) - 0.02336\ln(\delta B) - 0.00001589D - 6.671D/C \tag{13}$$

implying a final prediction for the total number of matches over a 10-year period of

$$\hat{N} = 10D\tau\hat{Y} = 10\lambda(\varphi(B) - \psi(B, C))\hat{Y}, \tag{14}$$

where we use the fact that $D = \lambda\varphi(B)$ (i.e., the number of patients who are unable to find a match in the BMR, using Equation (5)) and that $\tau = 1 - \psi(B, C)/\varphi(B)$, as given above. A plot of N against $\hat{N}$ is given in Figure 3b, which shows the accuracy of the fit across the full range of N.

Figure 3 (a) Declining CB matching probability over time and (b) Predicted vs. simulated UCB matches for regression model

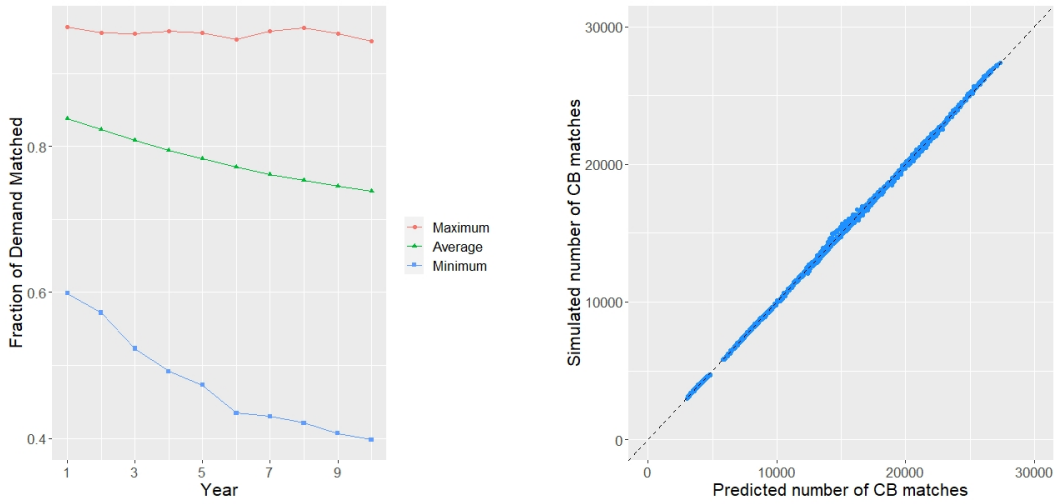

(a) (b)

In the next subsection, we will formulate and solve an approximate optimization model for the joint BMR/CBB design problem.

6.3. Joint optimization based on the predictive regression model

Now that we have the necessary parameters in hand and the issue of the changing CBB composition has been resolved, we can perform joint optimization of the BMR and the CBB. In particular, the modeling framework developed in Section 5.4 can be modified by making the following changes to Equation (11):

1. The expected number of CBB matches $\lambda(\varphi(B) - \psi(B,C))$ is replaced with $\lambda(\varphi(B) - \psi(B,C))\hat{Y}$, the average annual predicted CBB matches (using Equation (14)).
2. The value of a CBB match is reduced to $v_c - c_c$, since any CBU used must be replaced at a cost of c_c.
3. The annual cost of replacing the entire CBB (i.e., Cc_c) is replaced by the sum of the cost of replacing 5% of the CBB (the expired units) at the start of each year (i.e., $0.05Cc_c$), together with the annualized cost of establishing a CBB of size C and a BMR of size B in year one (i.e., $0.1(Cc_c + Bc_b)$).

Hence, the objective now is to maximize the function

$$H(B,C) = -0.1(Cc_c + Bc_b) + \begin{bmatrix} \lambda(1-\varphi(B))v_b + \lambda(\varphi(B) - \psi(B,C))\hat{Y}(v_c - c_c) \\ -0.05Cc_c - \alpha Bc_b - Ch \end{bmatrix} \tag{15}$$

or alternatively, to minimize

$$J(B,C) = \lambda v_b - H(B,C) = \begin{bmatrix} \lambda\varphi(B)(v_b - \hat{Y}(v_c - c_c)) + \lambda\psi(B,C)\hat{Y}(v_c - c_c) \\ +0.05Cc_c + \alpha Bc_b + Ch \end{bmatrix} + 0.1(Cc_c + Bc_b) \tag{16}$$

Figure 4 **3-D plot of** $J(B,C)$

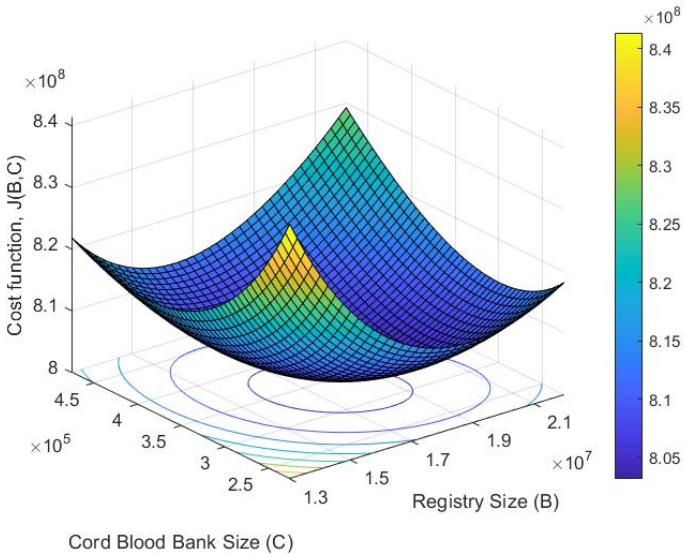

Since the convexity of $J(B,C)$ depends on the regression parameters, we were not able to prove that the function is jointly convex in general. However, over the range of values considered, we find that joint convexity always holds (see Figure 4). Since a closed form solution cannot be obtained, we numerically minimize $J(B,C)$. The global minimum of the function $J(B,C)$ is attained at $B^* \approx$ 17,490,000 and $C^* \approx$ 335,000. In comparison to the extreme policy, the optimal size of the BMR has decreased by approx. two million registered donors, while the optimal size of CBB has increased by a factor of five. The significantly larger optimal CBB size is expected because compared to the extreme policy in which all CBUs are discarded annually, the cost of maintaining a bank of size C where only 5% of CBUs are discarded annually is significantly less.

Under this optimal scenario, 12,065 of 13,000 patients find a match in the BMR, and 769 (i.e., 82.2% of the residual demand that spills over from the BMR) find a match in the CBB. This makes for a total of 12,834 matches, or a fill rate of 98.7%. The estimated annual costs for running the BMR and the CBB are $73.5 million and $51.8 million, respectively. Adding to this the 10-year amortized cost of expanding both institutions from their current sizes to the optimal ($62.9 million for the BMR and $14.1 million for the CBB), the total annual operating cost across the 10-year horizon of our study is therefore estimated at $202 million. This cost is 66% of that of the extreme policy, resulting in savings of more than $100 million per year, while an additional 16 lives are also saved annually. These numbers clearly demonstrate the advantages of operating an optimized joint BMR and CBB.

While these findings cannot be directly compared to those from other mathematical models, as ours is the first to assess both the BMR and CBB jointly, it is possible to compare with estimates using actual BMR/CBB composition data from the U.S. In particular, using composition data from the end of 2012,

Gragert et al. (2014) estimated the probability of finding a match to be 97% for white Europeans (who represent 75% of the demand for allogeneic HSCs (Health Resources and Services Administration 2020b)) and 95% for African Americans (15% of the demand). This was based on a BMR size of 10.5 million donors and a CBB size of 200,000 CBUs. Inputting into Equation (16) values of 10.5 million and 200,000 for the BMR and CBB sizes, respectively, we obtain a predicted matching rate of 97.6%, reassuringly close to Gragert et al. (2014).[11]

6.4. Sensitivity analysis

This study assumes 13,000 as the demand for allogeneic transplants. However, medical researchers predict that a wider variety of diseases may be treatable by HSC transplantation in the future (Liras 2010). Therefore, it is interesting to consider the impact of an increase in annual demand on system performance. To this end, in Table 2 we compare results when the BMR and CBB sizes are jointly optimized for $\lambda \in \{13000, 25000, 50000\}$. This shows that there are clear performance gains to be had as demand increases. First, the percentage of patients who find a match and the percentage of lives saved relative to the theoretical maximum increase. Second, after the optimum inventory levels have been reached, we can see that the average cost per life saved reduces significantly as demand increases, indicating economies of scale.

λ	B^* ($\times 10^6$)	C^* ($\times 10^5$)	# Matches		Fill Rate (%)	% Possible Lives Saved†	Operating Cost ($\times \$10^6$)‡	Cost per Life Saved ($\times \$10^3$)§
			BMR	CBB				
13,000	17.5	3.35	12,065	769	98.7	97.3	125	47
25,000	26.6	4.71	23,639	1,141	99.1	98.0	184	36
50,000	41.3	6.66	48,001	1,700	99.4	98.6	277	27

† Predicted lives saved divided by total potentially savable lives, calculated by assuming all λ patients receive a BM donation; ‡ Operating cost is the annual operating cost, excluding the amortized costs of expanding the BMR and CBB to their recommended levels; § Calculated by dividing the annual operating cost by the expected number of lives saved.

Table 2: Sensitivity Analysis on Demand

Another significant parameter that can influence policy recommendation is the VSL. We use a VSL of $6.5 million, as suggested by Bergstrom et al. (2009). However, the same authors acknowledge that estimates for a VSL can range from $4 million to $9 million. Therefore, we also investigate the sensitivity of our model to a change in the VSL. We find that with a VSL of $4 million, the optimal BMR and CBB

[11] In the same paper, the authors also report the probability of a match when excluding the CBB as a source (i.e., using only BM donors). Specifically, they report a 97% probability that a white European will find a match in the BMR and a 77% for African Americans. The average across other racial and ethnic groups is between 70 and 90%. Taking a weighted average of these probabilities, and using the midpoint of 80% for other groups, suggests that $\pm 92\%$ ($= [(97 * 0.75) + (77 * 0.15) + (80 * 0.1)]/100$) of patients should find a match in the BMR. With a BMR size of 10.5 million and a CBB size of 200k, our model reports that the expected number of patients who will find a match in the BMR is 11,711 (see Table 3). Similar to Gragert et al. (2014), our model predicts that 11,711/13,000 = 90.1% of patients would find a match in the BMR.

sizes reduce to 12.8 million and 250,000, respectively, with a corresponding fill rate of 98.1%. Interesting, this is very close to the current sizes of the BMR and CBB, meaning that a policy of maintaining current inventory levels is equivalent to valuing a life at approx. $4 million. Meanwhile, when using a VSL of $9 million, we find that 20.7 million registered donors in the BMR and 408,000 units of CB in the CBB is optimal, resulting in a fill rate of 99.0%. This follows intuitively from the fact that it is worth investing in more capacity to save additional lives when a life is valued more highly.

7. Conclusions, Implications and Future Scope

Public health services exist to provide citizens with access to life-saving and health-improving treatments and technologies. Supported by national and local governments, providers are tasked with allocating resources and services equitably and cost-effectively to improve quality of life and prevent unnecessary morbidity and mortality. That said, while many providers aim to be effective and efficient in order to meet this goal, more can always be done. It is estimated, for example, that one-third of medical spending in the U.S. is wasted due to excessive prices, unnecessary services, failures of care delivery and care coordination, administrative costs, and fraud and abuse (Berwick and Hackbarth 2012). In this paper, we focus on a specific health concern – the treatment of patients with hematological diseases such as leukemia via stem cell transplantation – to demonstrate how operations management tools combined with detailed context-specific knowledge can be applied to drive improvements in healthcare efficiency and, more importantly, save lives.

Specifically, we focus on the problem of allocating capacity between two alternative sources of HSCs, BMRs and CBBs, that have different operational and medical characteristics such that neither is preferred exclusively to the other. While this problem has been considered in the past, until now these two sources have been considered only in isolation, with the presence of the other source entirely ignored. Yet, we find that policy implications differ drastically when both sources are properly accounted for. For example, we find that if the U.S. were to rely only on the BMR as a source of HSCs, approx. 34 million registered donors would be required to ensure a sufficient supply of stem cells to minimize total annual costs. By comparison, when we account for UCB using an extreme policy under which all unused CBUs are discarded or donated at the end of each year, the recommended size of the BMR falls by 57% to just 19.7 million registered donors and nearly 308 additional patients can find a match each year. This finding highlights the significant value that CBUs offer as a flexible alternative source of HSCs.

While the extreme policy demonstrates the value of UCB, providing recommendations for the size of the BMR and CBB requires a mathematical model more closely matching the research context. As we discussed, this task is complicated by the fact that the quantity and distribution of demand spilling over to the CBB are highly contingent on the size of the BMR. Further, the distribution of supply in the CBB changes

over time and becomes increasingly misaligned with demand, resulting in fewer matches. Therefore, in tackling this problem, we first developed a simulation routine to capture these features. The results were then used to build a prediction model (achieving a fit of 99% as measured by the R^2 value) to identify the total number of lives saved over a 10-year period. Substituting the regression parameters into a more sophisticated mathematical model, we found approx. 17.5 million registered donors and 335 thousand CBUs in the BMR and CBB, respectively, to be the optimal policy. This allocation would cost $202 million per year while enabling 98.7% of patients to find a match. In comparing these numbers with an annual operating cost of $305 million (resp., $381 million) and a 97.8% (resp., 95.4%) match rate under the extreme policy (resp., BMR only policy), we can see the importance of accounting for inter-dependencies and temporal dynamics when jointly optimizing the sizes of the two institutions.

We can also use the joint optimization model specified in Equation (16) to estimate the number of matches, lives saved, and annualized costs under the status quo, i.e., by setting $B = 11.5$ million and $C = 265,000$ and leaving all other parameters unchanged. Doing so, we estimate the current potential for 11,780 BMR and 973 CBB matches per year, leaving 247 patients unmatched. In other words, under the optimal policy we are able to find matches for 33% $(= 1 - (13,000 - 12,834)/247)$ of the currently unmet demand. Moreover, since in the status quo more matches are found in the CBB and survival rates are lower with UCB HSCs than BM HSCs, 27 fewer lives are saved per year in the status quo than would be saved Meanwhile, the optimal policy results in 825 additional BMR matches $(= 12,605 - 11,780)$, each valued at $1.2 million, and 204 fewer CBB matches $(= 769 - 973)$, each valued at $903,000, resulting in an additional $806 million of value generated via matches each year. Subtracting the additional annual operating costs for the expanded BMR and CBB, the optimal policy provides nearly $770 million of additional value annually. After accounting for the costs of expanding the existing BMR and CBB to reach the optimal, over ten years we thus estimate a total increase in social surplus of approx. $6.9 billion following the adoption of the optimal policy.

A comparison of the findings under the different scenarios discussed above and in Section 6.3, using different values for the size of the BMR and CBB, is summarized in Table 3.

7.1. Policy implications

This paper assumes the existence of a central decision maker who makes capacity decisions. In reality, however, the organizational structure for unrelated stem cell transplants in the U.S. is highly decentralized, with BM donors and CBUs sourced from various NGOs and NPOs that separately maintain BMRs and operate public CBBs. The C.W. Bill Young Cell Transplantation Program (CWBYCTP), authorized by the U.S. Congress, motivates the relevance of the coordinated approach adopted in this paper (Resources and Administration 2018). In particular, the CWBYCTP is a program focused on increasing the number of BM

Setup	B ($\times 10^6$)	C ($\times 10^5$)	# Matches BMR	# Matches CBB	Fill Rate (%)	% Possible Lives Saved[†]	Op. Cost ($\times \$10^6$)[‡]	Annu. Cost ($\times \$10^6$)[§]
BM only	34	-	12,405	-	95.4	95.4	143	381
BM + CB (Extreme Policy)	19.7	0.66	12,134	579	97.8	96.7	219	305
BM + CB (Joint Model)	17.5	3.35	12,605	769	98.7	97.3	125	202
Bergstrom et al. (2009)	21	-	12,170	-	93.6	93.6	88	188
Gragert et al. (2014)	10.5	2	11,711	974	97.6	95.8	76	76
Current Status	11.5	2.65	11,780	973	98.1	96.3	90	90

[†] Predicted lives saved divided by total potentially savable lives, calculated by assuming all 13,000 patients receive a BM donation; [‡] Operating (Op.) cost is the annual operating cost, excluding the amortized costs of expanding the BMR and CBB to their recommended levels; [§] Annualized (Annu.) cost is the operating cost plus the 10-year amortized cost of expanding the BMR and CBB to their recommended levels.

Table 3: Summary of Results

and CB transplants for recipients matched to biologically unrelated donors. To achieve this, one of the primary objectives for the CWBYCTP is to develop a coordinated system for stem cell transplants from unrelated donors. Specifically, a single point of access, BM coordinating center, and CB coordinating center are combined to form a Single Point of Access-Coordinating Center (SPA-CC) contract.

In order to establish a center that effectively coordinates the functions described above, it is necessary from a policy perspective for a benchmark to exist towards which the SPA-CC can aspire and against which the contracted out party can be held accountable. Our model provides precisely this benchmark for a planning horizon of 10 years. Furthermore, the benchmark can also be used in order to devise incentives (e.g., subsidies) in order to encourage the different entities involved to move towards the first-best/coordinated solution. For example, the CWBYCTP funds a grant program for public CBBs, called the National Cord Blood Inventory (NCBI), which helps to fund the collection and maintenance of CBUs used in transplants. Meanwhile, the NMDP, the largest BMR with approx. 80% of the total registered donors in the U.S. (see Footnote 2), is federally funded. The results from our model can thus be directly applied for setting targets, devising incentives, and determining grants for these organizations.

It is important to also point out that although BM donor recruitment organizations operate as separate entities, the majority report directly to the Be The Match program. By setting targets for donor recruitment, Be The Match can therefore directly influence the number of donors on the BMR. Furthermore, although most public CBBs exist as separate not-for-profit entities, the majority of the demand for these banks arrives domestically via NMDP Cord or internationally via the Center for International Blood and Marrow Transplant Research (CIBMTR), both of which are managed by the Be The Match program. Furthermore, all public CBBs are regulated by the Health Resources and Services Administration (HRSA), who contract out the creation of the BMR and UCB registry to Be The Match. Given the importance of the HRSA and Be The Match in regulating public CBBs and driving demand to these institutions, they may be able to exert influence over the capacity of the individual banks.

7.2. Future research scope

Although this paper incorporates many of the key features of the setting in this policy recommendation, we also point to several possible extensions that might be productively considered in future research. First, an increase in microscope resolution has shifted the focus from three loci on chromosome pair six to four or even five loci. Hence, although phenotype information for most individuals in the BMR is only recorded for the six main alleles, the medical standard for BM (but not UCB) is shifting. Specifically, after a 6/6 match is found, further investigation will typically seek to determine whether the patients and donors match on additional alleles in other loci. Conditional on an initial 6/6 match, a 7/8 or 8/8 match is identified following further investigation in approx. 82% of cases, and the match will then typically be considered sufficient for transplantation (Howard et al. 2008, Gragert et al. 2014). While this advance is unlikely to have a major impact on our findings, as additional higher-resolution data is made available to researchers it could be incorporated into the model to refine the recommendations further.

Another recent medical development is progress towards transplanting two CBUs instead of one. With one limitation of UCB – especially for use in adults – being the low availability of units with a sufficiently high TNC content, double CBU transplants may remove a major barrier to their use. To date, outcomes of double CBU transplants have been found to be no better than those of standard CBU transplants (Wagner Jr et al. 2014); however, should outcomes improve, this feature could also be incorporated into our model.

In addition to such medical advances, the registries and banks are also becoming increasingly international and facilitating more cross-border matches. If rarer phenotype matches can be found by drawing from a larger, global pool of resources, this may reduce the need for local capacity. Evidence of scale economies in Section 6.4 also suggests that this trend may help in lowering costs. However, this shift may also give rise to incentive issues, e.g., countries may free ride on their partners' investments in capacity, especially if the phenotype distributions within the populations are similar. Investigating the mechanics and implications of these collaborations thus may provide further research opportunities.

7.3. Conclusion

Overall, our paper provides a dispassionate and rigorous model that can be used to inform policy. We present quantitative evidence to support ongoing calls – e.g., from the Be The Match program and the U.S. Department of Health and Human Services (Health Resources and Services Administration 2020b) – for more BM donors and CBUs in order to meet the demand for allogeneic transplants in the U.S. However, substantial resources and time are needed to ramp up capacity for both HSC sources. Therefore, beyond identifying the need for more capacity, establishing specific BMR and CBB capacity targets can help ensure that resources and time are allocated effectively to "get it right the first time." A rigorous analysis is thus vital for informing strategic planning and investment in these two institutions. As the first paper to describe and

solve this challenge as a joint optimization problem, our work thus allows for better informed, data-driven decision making and planning that can help to save lives while delivering significant social value.

Acknowledgments

References

Aly RM (2020) Current state of stem cell-based therapies: An overview. *Stem cell investigation* 7.

Anderson R, Ashlagi I, Gamarnik D, Rees M, Roth AE, Sönmez T, Ünver MU (2015) Kidney exchange and the alliance for paired donation: Operations research changes the way kidneys are transplanted. *Interfaces* 45(1):26–42.

Bassok Y, Anupindi R, Akella R (1999) Single-period multiproduct inventory models with substitution. *Operations Research* 47(4):632–642.

Be The Match (2020a) Joining the marrow registry: Members of the U.S. military. URL `https://bethematch.org/support-the-cause/donate-bone-marrow/join-the-marrow-registry/members-of-the-u-s--military/`, Accessed: 28 August 2020.

Be The Match (2020b) Matching patients with donors. URL `https://bethematch.org/transplant-basics/matching-patients-with-donors/`, Accessed: 24 August 2020.

Beatty PG (1994) The immunogenetics of bone marrow transplantation. *Transfusion Medicine Reviews* 8(1):45–58.

Beatty PG, Boucher KM, Mori M, Milford EL (2000) Probability of finding hla-mismatched related or unrelated marrow or cord blood donors. *Human Immunology* 61(8):834–840.

Bergstrom TC, Garratt RJ, Sheehan-Connor D (2009) One chance in a million: Altruism and the bone marrow registry. *American Economic Review* 99(4):1309–34.

Berwick DM, Hackbarth AD (2012) Eliminating waste in US health care. *JAMA* 307(14):1513–1516.

Besse KL, Preussler JM, Murphy EA, Denzen EM, Lill MC, Chell JW, Senneka MK, Majhail NS, Williams EP (2015) Estimating demand and unmet need for allogeneic hematopoietic cell transplantation in the united states using geographic information systems. *Journal of Oncology Practice* 11(2):e120–e130.

Blum A, Dickerson JP, Haghtalab N, Procaccia AD, Sandholm T, Sharma A (2020) Ignorance is almost bliss: Near-optimal stochastic matching with few queries. *Operations Research* 68(1):16–34.

Brown KS, Rao MS, Brown HL (2019) The future state of newborn stem cell banking. *Journal of clinical medicine* 8(1):117.

Dehn J, Spellman S, Hurley CK, Shaw BE, Barker JN, Burns LJ, Confer DL, Eapen M, Fernandez-Vina M, Hartzman R, et al. (2019) Selection of unrelated donors and cord blood units for hematopoietic cell transplantation: guidelines from the nmdp/cibmtr. *Blood* 134(12):924–934.

Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* 39(1):1–22.

Dickerson JP, Procaccia AD, Sandholm T (2019) Failure-aware kidney exchange. *Management Science* 65(4):1768–1791.

Ding Y, Ge D, He S, Ryan CT (2018) A nonasymptotic approach to analyzing kidney exchange graphs. *Operations Research* 66(4):918–935.

DKMS (2020) Registered Donors in the U.S. | DKMS. URL https://www.dkms.org/en, Accessed: 24 August 2020.

Fève F, Florens JP (2010) A mathematical model for bone marrow donors' registries and cord blood banks .

Giralt S, Bishop MR (2009) Principles and overview of allogeneic hematopoietic stem cell transplantation. *Hematopoietic Stem Cell Transplantation*, 1–21 (Springer).

Glorie KM, van de Klundert JJ, Wagelmans AP (2014) Kidney exchange with long chains: An efficient pricing algorithm for clearing barter exchanges with branch-and-price. *Manufacturing & Service Operations Management* 16(4):498–512.

Gökalp E, Gülpınar N, Doan XV (2020) Capacity planning for networks of stem-cell donation centers under uncertainty. *Production and Operations Management* 29(2):281–297.

Gragert L, Eapen M, Williams E, Freeman J, Spellman S, Baitty R, Hartzman R, Rizzo JD, Horowitz M, Confer D, et al. (2014) Hla match likelihoods for hematopoietic stem-cell grafts in the us registry. *New England Journal of Medicine* 371(4):339–348.

Harrison JM, Van Mieghem JA (1999) Multi-resource investment strategies: Operational hedging under demand uncertainty. *European Journal of Operational Research* 113(1):17–29.

Health Resources and Services Administration (2020a) Detailed Description of Donor Registry Data | Blood Stem Cell. URL https://bloodstemcell.hrsa.gov/data/donation-and-transplantation-statistics/donor-registry-data/detailed-description-donor#fig1, Last modified: April 2020; Accessed: 24 August 2020.

Health Resources and Services Administration (2020b) Frequently Asked Questions | Blood Stem Cell. URL https://bloodstemcell.hrsa.gov/about/faqs, Accessed: 22 August 2020.

Health Resources and Services Administration (2020c) The Need for More Marrow Donors | Blood Stem Cell. URL https://bloodstemcell.hrsa.gov/donor-information/donate-bone-marrow/need-more-marrow-donors, Accessed: 24 August 2020.

Health Resources and Services Administration (2020d) Transplant activity report. URL https://bloodstemcell.hrsa.gov/data/donation-and-transplantation-statistics/transplant-activity-report, Updated on: October 2020; Accessed via the Internet Archive on: 5 March 2021.

Howard DH, Meltzer D, Kollman C, Maiers M, Logan B, Gragert L, Setterholm M, Horowitz MM (2008) Use of cost-effectiveness analysis to determine inventory size for a national cord blood bank. *Medical Decision Making* 28(2):243–253.

Kekre N, Antin JH (2014) Hematopoietic stem cell transplantation donor sources in the 21st century: choosing the ideal donor when a perfect match does not exist. *Blood* 124(3):334–343.

Kollman C, Abella E, Baitty RL, Beatty PG, Chakraborty R, Christiansen CL, Hartzman RJ, Hurley CK, Milford E, Nyman JA, et al. (2004) Assessment of optimal size and composition of the us national registry of hematopoietic stem cell donors. *Transplantation* 78(1):89–95.

Lee J, Gaur V, Muthulingam S, Swisher GF (2016) Stockout-based substitution and inventory planning in textbook retailing. *Manufacturing & Service Operations Management* 18(1):104–121.

Liras A (2010) Future research and therapeutic applications of human stem cells: general, regulatory, and bioethical aspects. *Journal of translational medicine* 8(1):131.

Majhail NS, Mothukuri JM, Brunstein CG, Weisdorf DJ (2009) Costs of hematopoietic cell transplantation: comparison of umbilical cord blood and matched related donor transplantation and the impact of posttransplant complications. *Biology of Blood and Marrow Transplantation* 15(5):564–573.

Mori M, Beatty PG, Graves M, Boucher KM, Milford EL (1997) Hla gene and haplotype frequencies in the north american population: The national marrow donor program donor registry1. *Transplantation* 64(7):1017–1027.

Navarrete C, Contreras M (2009) Cord blood banking: a historical perspective. *British Journal of Haematology* 147(2):236–245.

Parent's Guide to Cord Blood Foundation (2020) Public cord blood banking in United States. URL `https://parentsguidecordblood.org/en/public-banking/united-states`, Accessed: 26 August 2020.

Resources H, Administration S (2018) Annual Progress Report on the C.W. Bill Young Cell Transplantation Program and National Cord Blood Inventory Program. URL `https://bloodstemcell.hrsa.gov/sites/default/files/bloodstemcell/about/legislation/fiscal-year-2018-annual-progress-report.pdf`, Accessed: 1 March 2020.

Roth AE, Sönmez T, Ünver MU (2004) Kidney exchange. *The Quarterly Journal of Economics* 119(2):457–488.

Roth AE, Sönmez T, Ünver MU (2005) Pairwise kidney exchange. *Journal of Economic theory* 125(2):151–188.

Schlapp J, Fleischmann M (2018) Multiproduct inventory management under customer substitution and capacity restrictions. *Operations Research* 66(3):740–747.

Song JS, van Houtum GJ, Van Mieghem JA (2020) Capacity and inventory management: Review, trends, and projections. *Manufacturing & Service Operations Management* 22(1):36–46.

Sonnenberg FA, Eckman MH, Pauker SG (1989) Bone marrow donor registries: the relation between registry size and probability of finding complete and partial matches [see comments]. *Blood* 74(7):2569–2578.

Su X, Zenios SA (2006) Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Science* 52(11):1647–1660.

Van Mieghem JA (1998) Investment strategies for flexible resources. *Management Science* 44(8):1071–1078.

Van Mieghem JA (2003) Commissioned paper: Capacity management, investment, and hedging: Review and recent developments. *Manufacturing & Service Operations Management* 5(4):269–302.

Wagner Jr JE, Eapen M, Carter S, Wang Y, Schultz KR, Wall DA, Bunin N, Delaney C, Haut P, Margolis D, et al. (2014) One-unit versus two-unit cord-blood transplantation for hematologic cancers. *New England Journal of Medicine* 371(18):1685–1694.

World Marrow Donor Association (2020) Total number of donors and cord blood units. URL `https://statistics.wmda.info/`, Last modified: 23 August 2020; Accessed: 24 August 2020.

Zenios SA (2002) Optimal control of a paired-kidney exchange program. *Management Science* 48(3):328–342.

E-companion to:
"Be the Match: Optimizing Capacity Allocation for Allogeneic Stem Cell Transplantation"

EC.1. Proofs

Proof of Proposition 1: First we will find the probability of exactly a 4/6 match. Note that for phenotype $AaBbDd$ the probability of exactly a 4/6 match with mismatches in A and B is given by

$$p_{.a.bDd} - p_{.aBbDd} - p_{Aa.bDd} + p_{AaBbDd}.$$

Similar expressions apply to the probabilities of exactly a 4/6 match with mismatches in A and b, A and D, A and d, a and B, $\cdots$, b and d. The sum over all 12 pairs then gives us the probability of exactly a 4/6 match for $AaBbDd$. Hence we have

$$\begin{aligned}
m_4^{AaBbDd} &= m_5^{AaBbDd} + p_{.a.bDd} - p_{.aBbDd} - p_{Aa.bDd} + p_{AaBbDd} \\
&\quad + p_{.aB.Dd} - p_{.aBbDd} - p_{AaB.Dd} + p_{AaBbDd} + \cdots \\
&\quad + p_{AaB.D.} - p_{AaB.Dd} - p_{AaBbD.} + p_{AaBbDd} \\
&= m_5^{AaBbDd} + 12p_{AaBbDd} \\
&\quad + p_{.a.bDd} + p_{.aB.Dd} + p_{.aBb.d} + p_{.aBbD.} + p_{A..bDd} + p_{A.B.Dd} \\
&\quad + p_{A.Bb.d} + p_{A.BbD.} + p_{Aa.b.d} + p_{Aa.bD.} + p_{AaB..d} + p_{AaB.D.} \\
&\quad - 4(p_{.aBbDd} + p_{A.BbDd} + p_{Aa.bDd} + p_{AaB.Dd} + p_{AaBb.d} + p_{AaBbD.})
\end{aligned}$$

Using Equation (2) and some more rearranging of terms then gives Equation (3).

Proof of Proposition 2: It is sufficient to show that $\varphi(B)$ as defined in Equation (5) is convex. But this is immediate since $x^B = e^{B\ln x}$ is convex in B for every $x > 0$, so all the terms $p_i(1-\delta p_i)^B = p_i e^{B\ln(1-\delta p_i)}$ in the summation in the right-hand side of Equation (5) are convex in B.

Proof of Proposition 3: The first term inside the square braces on the right hand side in Equation (12) is jointly convex in B and C since $\varphi(B)$ is convex in B (as shown in the proof of Proposition 2) and $v_b - v_c > 0$. All the terms except $\psi(B,C)$ are linear in B and C. So it is sufficient to show that $\psi(B,C)$ is jointly convex in B and C. But $\psi(B,C)$ is a sum of functions of the form $e^{xB+yC+z}$, each of which are jointly convex in B and C as the composition of an increasing convex function and a convex function. So $\psi(B,C)$ is the sum of jointly convex functions and therefore jointly convex itself.

EC.2. Procedure for Simulation

The simulation algorithm described in Section 6.1 of the main paper is documented here in Algorithm 1.

Two important features of the simulation algorithm require further elaboration. First, we implement a priority rule in which 6/6 matches are always chosen first (if available), followed by 5/6 and 4/6 matches. This is due to faster

patient recovery times when matching on more alleles, as noted in Section 2.3. Second, in the case of multiple matches, one might prefer to allocate the CBU with a higher TNC content to the patient, since lower TNC content has been associated with higher mortality (Dehn et al. 2019). However, there are several factors aside from the TNC content that determines whether a physician will choose a particular CBU. One major factor is the time to access the CBU. As compromises will need to be made in reality, we instead allocate randomly.

Algorithm 1: Procedure for Simulation

Inputs: Scalars B, C, δ, D and vectors p, q, distributions of η and ζ;

Initialize:

Draw CBU sample of size C (phenotypes from p and TNC content levels $\sim \eta$);

Counters: nmatches = 0;

Repeat 10 times:

Draw patient sample of size D (phenotypes from q and minimum TNC content level $\sim \zeta$);

For Each patient **In** patient sample **Do**

1. Find all 6/6 matches with sufficient TNC content in CBB. If none exist, go to 2, else randomly select one, remove it from the CBB, nmatches++; break;
2. Find all 5/6 matches in CBB. If no 5/6 match with sufficient TNC content exists, go to 3, else randomly select one matching unit with sufficient TNC content from the phenotype with the most 5/6 matches and at least one match with sufficient TNC content, remove it from the CBB; nmatches++; break;
3. Find all 4/6 matches in CBB. If no 4/6 match exists with sufficient TNC content, go to **Next**, else randomly select one matching unit with sufficient TNC content from the phenotype with the most 4/6 matches and at least one match with sufficient TNC content, remove it from the CBB; nmatches++; break;

Next

Randomly remove $0.05C$ units from the CB registry (expiry) and replace from the population distribution p to fill up to size C;

End Repeat

Report Results: Output nmatches

EC.3. Medical comparison of stem cells from bone marrow versus stem cells from cord blood

As noted in Section 2.3, in addition to differences in their operational characteristics, the two sources of HSCs differ along certain medical dimensions that are important for our analysis. We describe these in more detail below.

Total nucleated cell count. As discussed in Section 2.3, the total nucleated cell (TNC) content content of CBUs is typically much lower than that of BM. In particular, the TNC content of most CBUs ranges from 50×10^7 to 400×10^7, while a dose of at least 2.5×10^7 per kilogram of body weight is required for successful transplantation. A TNC content in the range of 180-230$\times 10^7$ cells is therefore typically needed for an average-sized adult in the U.S., yet only 15% of CBUs in public CBBs have counts above 175$\times 10^7$ (Kapinos et al. 2017).

Graft-versus-host disease. The most frequent complication that occurs after an allogeneic transplant is called graft-versus-host disease (GvHD), a complication of transplantation in which the donor cells attack the cells of the recipient. Between 30% and 70% of patients with a donor stem cell transplant get some form of GvHD, which can vary in severity from mild to life-threatening. GvHD can be either acute, typically coming on 10 to 90 days after the transplant and lasting on average for 25 days, or chronic, starting and lasting for anywhere between 90 and 600 days after surgery (American Cancer Society 2020). Overall, lower rates of GvHD and fewer complications have been found in UCB transplantation patients than in BM transplantation patients (Chen et al. 2013). However, while GvHD can have serious and unpleasant symptoms, in most cases it is not lethal and symptoms can be managed with immunosuppressive drugs and steroids. Therefore, this condition does not often provide a sufficiently compelling reason to prefer UCB over BM as a source of HSCs.

Graft failure. A more serious complication that can occur following stem cell transplantation is known as graft failure, which occurs when engraftment does not succeed and donor cells then fail to integrate properly into the recipient's cell population. This can either be due to a lack of initial engraftment of donor cells (primary graft failure) or due to the loss of donor cells after initial engraftment (secondary graft failure). Although graft failure is not a major concern in BM transplants, the additional incidence rate of graft failure among UCB recipients relative to BM recipients is estimated at approximately five percentage points (Brunstein et al. 2010). As compared to GvHD, graft failure can be more lethal. Consequently, CBUs tend to work best when administered to children and young adults since these patients are still in their formative years and their bodies can better adapt to and accommodate the CBUs, rendering the issue of graft failure a lesser concern (Sanz and Sanz 2002).

Slow engraftment. Another risk factor associated with HSC transplant in general is slow engraftment. The engraftment rate measures the time taken for the BM/CB cells to proliferate into the patient's bone marrow. This problem is especially a concern when it comes to mismatched HLA (i.e., 5/6 and 4/6) CB transplants. Although there is no evidence of a difference in overall mortality rates for HLA mismatched CB transplants, the engraftment rate is slower when there is a HLA mismatch, which can increase the recovery time (Ruggeri et al. 2014).

EC.4. Glossary of Terms and Definitions

EC.4.1. Acronyms

BM	Bone marrow
BMR	Bone marrow registry
CB	Cord blood (cells)
CBB	Cord blood bank
CBU	Cord blood unit
CIBMTR	Center for International Blood and Marrow Transplant Research
CWBYCTP	C.W. Bill Young Cell Transplantation Program
HLA	Human leukocyte antigens
HRSA	Health Resources and Services Administration
HSC	Hematopoietic stem cell
MUD	Matched unrelated (bone marrow) donor
NCBI	National Cord Blood Inventory (Program)
NMDP	National marrow donor program
TNC	Total nucleated cell
UCB	Umbilical cord blood

EC.4.2. Medical Terms and Definitions

The medical definitions provided in the section are referenced from Kenyon and Babic (2018) and the Anthony Nolan Foundation (2021).

Alleles (gene): A gene essentially is a part (locus) of the DNA structure and specific genes in the DNA decide specific traits of an individuals. Usually, the genes occur in pairs. Alleles are the different variations of the same gene found in a species. In the context of our study, HLA-A, HLA-B, and HLA-DRB1 are the three genes (basically protein structure) responsible for the immunity of an individual. Human beings can have different proteins (variation) on these three genes and each variation is an allele.

Allogeneic transplants: The stem cells in allogeneic transplants are from a source other than (outside) the patient. This may be a BM donor or a CBU. The transplant can either be matched or mismatched with the source.

Autologous transplants: The stem cells in an autologous transplants come from the same person who will get the transplant. In other words, the patient is their own donor.

Cord Blood Units (CBUs): CBUs are a source of HSCs extracted from the umbilical cord blood (UCB) (placenta) that is collected shortly after childbirth and processed to be stored in cryogenic facilities at cord blood banks for future use.

Engraftment: Engraftment is the process by which the transplanted hematopoietic stem cells (HSC) from the source make their way (homing) to the free bone marrow of the host where they can find optimal conditions to survive and proliferate.

Graft Failure: Graft failure is defined as the lack of hematopoietic cell engraftment (i.e., lack of survival and adequate proliferation) following autologous or allogeneic stem cell transplant.

Graft Rejection: The term graft rejection refers to the immune-mediated rejection of the donor cells by the residual host cells because of genetic disparity between the recipient and the donor. Therefore, this term is only relevant to allogeneic transplants.

Graft versus Host Disease (GvHD): GvHD is the condition in which the graft (the transplanted bone marrow cells accepted by the host) attacks and injures specific tissues of the recipient (i.e., the opposite of graft rejection). This can lead to localized medical conditions that are mild to life-threatening in severity.

Haplotype: The HLA in each chromosome is inherited as a "set" of the three HLA groups: A, B, DR from each parent. This set is known as a haplotype. Thus, the combination of alleles for HLA-A, HLA-B, and HLA-DRB1 on each individual chromosome is known as a haplotype.

Hematopoietic stem cell (HSC): Stem cells, in general, are undifferentiated or partially differentiated cells that can differentiate into various types of cells and proliferate indefinitely to produce more of the same stem cell. Hematopoietic stem cells (HSCs) are the stem cells that are found in the bone marrow and give rise to other blood cells through a process called haematopoiesis. Several medical conditions and treatments can destroy the HSCs and may necessitate replacement in the form of transplant from either an internal or external source.

Histocompatibility: Histocompatibility is the state in which a donor and recipient share antigens (i.e., HLAs) so that a graft is accepted and remains functional. In the context of our study, a histocompatibility match between a BM donor or CBU and a recipient is determined by the degree of similarity between the alleles that make up their phenotypes. The (histocompatible) match is classified as a 6/6 or 5/6 or 4/6 match based on the number matches in the allele of the phenotype between the recipient and the donor/CBU.

Human Leukocyte Antigens (HLAs): In general, antigens are proteins on the surface of the cells in the body. Specifically, Human Leukocyte Antigens (HLAs) are antigens found in chromosome pair 6 (one of the 23 pairs of chromosomes). A chromosome pair is made up of two individual chromosomes, one inherited from each parent. There are many different HLAs in the chromosome pair 6, but the ones that are most relevant for the transplantation are HLA-A, HLA-B, and HLA-DRB1.

Matched Unrelated Donor (MUD): These are individual voluntary bone marrow donors with a 6/6 phenotype match with the patient. The HSCs required for transplant are extracted from the bone marrow tissue/blood stream of the BM donors.

Matched transplant: For a BM donor/CBU and a receiver pair, if the genetic codes at all six alleles ("$AaBbDd$") of the phenotype perfectly matches one-to-one, then it is a (fully) matched transplant. It is also referred to as a 6/6 match. It should be noted that only 6/6 matches are feasible for BM transplant to proceed successfully.

Mismatched transplant: For a BM donor/CBU and a receiver pair, if there is a mismatch in the genetic code in at least one of the six alleles, then it is referred to as a mismatched transplant. In the context of our study, mismatched transplants are feasible only for CB transplants where up to 2 mismatches are permissible. This leads to usage of the terms 5/6 match (1 mismatched) and 4/6 match (2 mismatched) in the context of CB transplants.

Phenotype: A phenotype is the combined array of HLAs found in the pair of chromosomes (i.e., taking both haplotypes together). The phenotype is typically represented by a string of length six, "$AaBbDd$", where Aa represents the two copies of genetic code (i.e., the two alleles, one from each haplotype) at locus HLA-A, and with Bb and Dd representing the two alleles at HLA-B and HLA-DRB1, respectively.

Total nucleated cell (TNC) content: The Total Nucleated Cell (TNC) count is the numerical measure of the hematopoietic stem cell (HSC) count after umbilical cord blood (UCB) is processed to extract the HSCs. It is a direct measure of the quality of the extracted/isolated stem cell from the UCB. The TNC content of most CBUs ranges from 50×10^7 to 400×10^7, while a dose of at least 2.5×10^7 per kilogram of recipient's body weight is required for successful transplantation.

EC.5. Discounted Model

In this section, we investigate how discounting can be incorporated and how it impacts the optimal registry and cord blood bank size we found in Section 6.3. In the health economics literature, it is commonly assumed that both the the costs and the health benefits have the same discount rate. Accordingly, we discount the total value of matches every

year and the annual operating cost while the cost of expanding the BMR and the CBB are kept as upfront costs. So we can modify Equation (15) to reflect the above mentioned changes. Specifically, we want to maximize

$$\tilde{H}(B,C) = -0.1(Cc_c + Bc_b) + 0.1\sum_{t=1}^{10}\beta^t \begin{bmatrix} \lambda(1-\varphi(B))v_b + \lambda(\varphi(B)-\psi(B,C))\hat{Y}_t(v_c - c_c) \\ -0.05Cc_c - \alpha Bc_b - Ch \end{bmatrix}, \tag{EC.1}$$

where $(\varphi(B)-\psi(B,C))\hat{Y}_t$ is the estimated probability that a demand in year t is filled from the CBB. It remains to specify how the estimates for $\hat{Y}_t$ are obtained.

From the simulation in Section 6.1, we have the number of matches n_{tj} in each year of the simulation ($t = 1,\ldots,10$) and for each simulation run ($j = 1,\ldots 2520$). Define $s_{tj} = \frac{n_{tj}}{\sum_{u=1}^{10} n_{uj}}$ and $S_{tj} = \sum_{u=1}^{t} s_{uj}$. Thus s_{tj} is the fraction of all CBB matches in simulation run j that occurred in year t, and S_{jt} is the fraction of all CBB matches in simulation run j that occurred in years 1 through t. We use a regression model to predict the cumulative shares of matches. Hence we want to minimize

$$\min_{\beta} \sum_{j=1}^{2520}\sum_{t=1}^{10}(L(\beta,t) - S_{tj})^2 \tag{EC.2}$$

for some suitably chosen function $L(\beta,t)$. We use the function $L(\beta,t) = t + \beta_1 t(1-t) + \beta_2 t(1-t)(1+t)$ Note that $L(\beta,0) = 0$ and $L(\beta,1) = 1$, which are the known points in the data. The least squares fit is obtained when $\hat{\beta}_1 = 0.11790097$ and $\hat{\beta}_2 = -0.0271953$. Hence our estimate for $\hat{Y}_t$ is as follows:

$$\hat{Y}_t = 10\hat{Y}(L(\hat{\beta},t) - L(\hat{\beta},t-1)) \tag{EC.3}$$

where $\hat{Y}$ is the average efficiency of CBB over 10 years (see Section 6.3).

Figure EC.1 shows side by side box-and-whisker plots of the distributions of $\{s_{jt} : j = 1,\ldots,2520\}$ as well as the estimates $\ell_t = L(\hat{\beta},t) - L(\hat{\beta},t-1)$ for $t = 1,\ldots,10$. The predicted values virtually coincide with the averages. The plot clearly shows the declining efficiency of the CBB over time.

Figure EC.1 Average and estimated efficiencies over time.

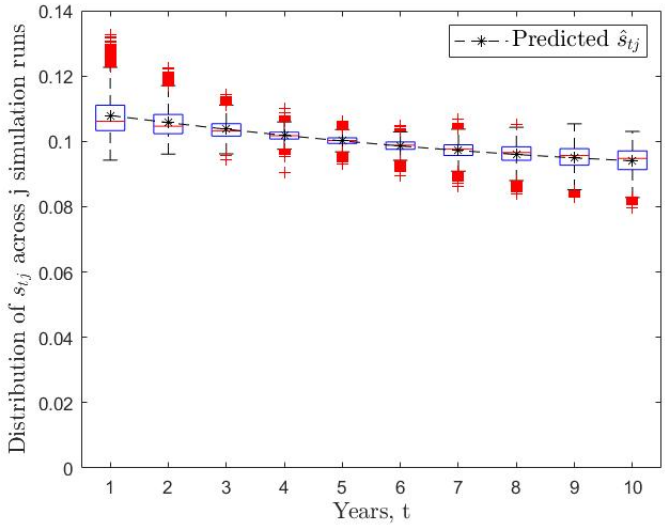

In Table EC.1 we show the results of jointly maximizing Equation (EC.1) over BMR size (B) and CBB size (C) for discount rates ranging from 0% to 10%. The optimal BMR and CBB sizes for 0% give the optimal sizes reported in Section 6.3. To isolate the impact of the declining values of $\hat{Y}_t$ from the impact of discounting, we compared the optimal BMR and CBB sizes in Table EC.1 with those obtained by maximizing

$$\bar{H}(B,C) = -0.1Cc_c - 0.1Bc_b + 0.1\bar{\beta}\begin{bmatrix} -0.05Cc_c - \alpha Bc_b + \lambda(1-\varphi(B))v_b \\ +\lambda(\varphi(B) - \psi(B,C))\hat{Y}(v_c - c_c) - Ch \end{bmatrix} \tag{EC.4}$$

where $\bar{\beta} = 0.1\sum_{t=1}^{10}\beta^t$, which corresponds to multiplying all of the cost parameters in the square brackets of Equation (EC.4) by $\bar{\beta}$. We don't report the results, but in no case is the difference more than 1000 in BMR size or 250 in CBB size. We conclude that the impact of discounting on optimal BMR ad CBB sizes can be adequately assessed by simple adjustments to the parameters c_c, c_b, v_b, v_c and h in the term within square brackets in $H(B,C)$ (see (15)), if desired.

Discount Rate (%)	0	1	2	3	4	5	6	7	8	9	10
Optimal BM Size (B^*) $\times 10^6$	17.5	17.1	16.6	16.2	15.8	15.4	15.1	14.7	14.4	14.1	13.8
Optimal CB Size (C^*) $\times 10^3$	335	331	328	324	321	317	314	311	307	304	300

.

Table EC.1: Optimal BM registry (B^*) and CB bank size (C^*) at different discount rates

References for e-Companion

American Cancer Society (2020) Stem Cell or Bone Marrow Transplant Side Effects. URL `https://www.cancer.org/treatment/treatments-and-side-effects/treatment-types/stem-cell-transplant/transplant-side-effects.html`, Last modified: 20 March 2020; Accessed: 24 August 2020.

Anthony Nolan Foundation (2021) Histocompatibility and immunogenetics. URL `https://www.anthonynolan.org/clinicians-and-researchers/histocompatibility-and-immunogenetics`, Updated on: 21 February 2021; Accessed via the Internet Archive on: 5 March 2021.

Brunstein CG, Gutman JA, Weisdorf DJ, Woolfrey AE, DeFor TE, Gooley TA, Verneris MR, Appelbaum FR, Wagner JE, Delaney C (2010) Allogeneic hematopoietic cell transplantation for hematologic malignancy: relative risks and benefits of double umbilical cord blood. *Blood* 116(22):4693–4699.

Chen Yh, XU Lp, Liu Dh, Huan C, Zhang Xh, Wei H, Wang Fr, Wang Jz, Yu W, Huang Xj, et al. (2013) Comparative outcomes between cord blood transplantation and bone marrow or peripheral blood stem cell transplantation from unrelated donors in patients with hematologic malignancies: a single-institute analysis. *Chinese Medical Journal* 126(13):2499–2503.

Dehn J, Spellman S, Hurley CK, Shaw BE, Barker JN, Burns LJ, Confer DL, Eapen M, Fernandez-Vina M, Hartzman R, et al. (2019) Selection of unrelated donors and cord blood units for hematopoietic cell transplantation: guidelines from the nmdp/cibmtr. *Blood* 134(12):924–934.

Kapinos KA, Briscombe B, Gracner T, Strong A, Whaley C, Hoch E, Hlavka JP, Case SR, Chen PG (2017) Public cord blood banks: Worthy of national investment. Technical report, RAND Corporation, URL `https://www.rand.org/pubs/research_briefs/RB9977.html`.

Kenyon M, Babic A (2018) *The European Blood and Marrow Transplantation Textbook for Nurses: Under the Auspices of EBMT* (Springer Nature).

Ruggeri A, Labopin M, Sormani MP, Sanz G, Sanz J, Volt F, Michel G, Locatelli F, De Heredia CD, O'Brien T, et al. (2014) Engraftment kinetics and graft failure after single umbilical cord blood transplantation using a myeloablative conditioning regimen. *haematologica* 99(9):1509.

Sanz M, Sanz G (2002) Unrelated donor umbilical cord blood transplantation in adults. *Leukemia* 16(10):1984–1991.