Stochastic Systems
2013, Vol. 3, No. 1, 38–95
DOI: 10.1214/12-SSY074

ON THE CONVERGENCE OF SIMULATION-BASED ITERATIVE METHODS FOR SOLVING SINGULAR LINEAR SYSTEMS*

By Mengdi Wang and Dimitri P. Bertsekas

Massachusetts Institute of Technology

We consider the simulation-based solution of linear systems of equations, $Ax = b$, of various types frequently arising in large-scale applications, where A is singular. We show that the convergence properties of iterative solution methods are frequently lost when they are implemented with simulation (e.g., using sample average approximation), as is often done in important classes of large-scale problems. We focus on special cases of algorithms for singular systems, including some arising in least squares problems and approximate dynamic programming, where convergence of the residual sequence $\{Ax_k - b\}$ may be obtained, while the sequence of iterates $\{x_k\}$ may diverge. For some of these special cases, under additional assumptions, we show that the iterate sequence is guaranteed to converge. For situations where the iterates diverge but the residuals converge to zero, we propose schemes for extracting from the divergent sequence another sequence that converges to a solution of $Ax = b$.

1. Introduction. We consider the solution of the linear system of equations

$$Ax = b,$$

where A is an $n \times n$ real matrix and b is a vector in $\Re^n$, by using approximations of A and b, generated by simulation. We assume throughout that the system is consistent, i.e., it has at least one solution. We consider iterative methods of the general form

$$x_{k+1} = x_k - \gamma G(Ax_k - b), \tag{1}$$

where γ is a positive stepsize, and their variants, where in place of A and b, we use simulation-generated approximations A_k, b_k, G_k, with $A_k \to A$, $b_k \to b$, $G_k \to G$:

$$x_{k+1} = x_k - \gamma G_k(A_k x_k - b_k). \tag{2}$$

Received May 2012.

*This work is supported by the Air Force Grant FA9550-10-1-0412 and by the LANL Information of Science and Technology Institute Grant 67870-001-08.

AMS 2000 subject classifications: Primary 15A06; secondary 60H99, 65C05.

Keywords and phrases: Stochastic algorithm, singular system, Monte-Carlo estimation, simulation, proximal method, regularization, approximate dynamic programming.

Most stationary iterative methods for solving the system are of the form (1), including projection, proximal, and splitting algorithms, as described in many books on iterative methods (see the references cited in Section 2.2). The choice of the stepsize γ in the simulation-based iteration (2) is the same as in the deterministic iteration (1). In some methods, such as proximal and splitting, a value $\gamma = 1$ is permissible and guarantees convergence, while in other methods such as projection, the proper value of γ for convergence may need to be estimated. Alternatively, one may consider using a diminishing stepsize sequence $\{\gamma_k\}$ with $\sum_{k=0}^{\infty} \gamma_k = \infty$; most of our convergence analysis applies to this case as well (see Section 2.4).

In our related paper [WB11] we showed that when A is singular, methods of the form (1) that are convergent, may easily become divergent when the entries of A and b are corrupted by simulation noise as in Eq. (2). We then introduced a general stabilization mechanism into iteration (2) that restores its convergence. The idea there was to shift the eigenvalues of $I-\gamma GA$ by a negative amount $-\delta_k$ into the unit circle, and then gradually reduce δ_k to 0, but at a rate that is slow enough to suppress the effects of the simulation noise $(G_kA_k - GA, b_k - b)$. In fact an entire class of stabilization schemes was proposed and analyzed in [WB11]. For example, a special stabilization scheme for proximal iterations was given, which shifts instead the eigenvalues of A.

In this paper, we discuss two special cases of systems and algorithms of the form (2), which yield a solution of $Ax = b$ without a stabilization mechanism. One such special case is when iteration (2) is *nullspace consistent* in the sense that the nullspaces of A_k and G_kA_k coincide with the nullspace of A. This case arises often in large-scale least squares problems and dynamic programming applications, as we will discuss in Section 2.2. The second case arises when the original system is transformed to the symmetric system $A'\Sigma^{-1}Ax = A'\Sigma^{-1}b$, and a *proximal algorithm* that uses quadratic regularization is applied to the transformed system with a stepsize $\gamma = 1$.

In both of these special cases, the sequence of residuals $\{Ax_k - b\}$ generated by iteration (2) typically converges to 0, but the sequence of iterates $\{x_k\}$ may diverge. To address this situation, we provide algorithms that extract from $\{x_k\}$ another sequence $\{\hat{x}_k\}$ that converges to a solution of $Ax = b$.

The paper is organized as follows. In Section 2 we summarize the convergence analysis for the deterministic iteration (1), including a necessary and sufficient condition for its convergence and a decomposition of the iteration that provides the basis for analysis of its simulation-based counterpart (2). Then we introduce simulation-based variants of the iterative algorithms, illustrate a few applications to practical large linear systems, and discuss the

related convergence issues, including the choice/estimation of the stepsize γ in the presence of stochastic noise. In Sections 3 and 4 we discuss the convergence properties of iteration (2) for the special cases we noted earlier. In Section 5 we give examples and prove divergence of the iterates or even the residuals under various conditions. In Section 6 we discuss how to estimate a matrix of projection onto the nullspace of A, which can be used to extract from $\{x_k\}$ another sequence that converges to a solution. Finally, in Section 7 we present computational results that support our analysis.

We summarize our terminology, our notation, and some basic facts as follows. A vector x is viewed as a column vector, while x' denotes the corresponding row vector. The standard Euclidean norm of a vector x is $\|x\| = \sqrt{x'x}$. For a matrix M, we use M' to denote its transpose. The nullspace and range of a matrix M are denoted by $\mathbf{N}(M)$ and $\mathbf{R}(M)$, respectively. We use $M^\dagger$ to denote the Moore-Penrose pseudoinverse of M (see the book by Ben-Israel and Greville [BIG74], among other sources, for discussions of its properties). We will use the fact that $M(M'M)^\dagger M'$ is the operator of orthogonal projection on $\mathbf{R}(M)$. For two square matrices A and B, the notation $A \sim B$ indicates that A and B are related by a similarity transformation and therefore have the same eigenvalues. We denote by $\rho(A)$ the spectral radius of A, and we denote by $\|A\|$ the Euclidean matrix norm of a matrix A, so that $\|A\|$ is the square root of the largest eigenvalue of $A'A$. We have $\rho(A) \leq \|A\|$, and we will use the fact that if $\rho(A) < 1$, there exists a weighted norm $\|\cdot\|_P$, defined using an invertible matrix P as $\|x\|_P = \|P^{-1}x\|$ for any $x \in \Re^n$, such that the corresponding induced matrix norm $\|A\|_P = \max_{\|x\|_P=1} \|Ax\|_P$ satisfies $\|A\|_P < 1$ (see the book by Stewart [Ste73], Th. 3.8).

If A and B are real square matrices, we write $A \preceq B$ or $B \succeq A$ to denote that the matrix $B - A$ is positive semidefinite, i.e., $x'(B - A)x \geq 0$ for all x. Similarly, we write $A \prec B$ or $B \succ A$ to denote that the matrix $B - A$ is positive definite, i.e., $x'(B-A)x > 0$ for all $x \neq 0$. We have $A \succ 0$ if and only if $A \succ cI$ for some positive scalar c [take c in the interval $\big(0, \min_{\|x\|=1} x'Ax\big)$].

If $A \succ 0$, the eigenvalues of A have positive real parts (see Theorem 3.3.9, and Note 3.13.6 of Cottle, Pang, and Stone [CPS92]). Similarly, if $A \succeq 0$, the eigenvalues of A have nonnegative real parts (since if A had an eigenvalue with negative real part, then for sufficiently small $\delta > 0$, the same would be true for the positive definite matrix $A + \delta I$ - a contradiction). For a singular matrix A, the algebraic multiplicity of the 0 eigenvalue is the number of 0 eigenvalues of A. This number is greater or equal to the dimension of $\mathbf{N}(A)$ (the geometric multiplicity of the 0 eigenvalue, i.e., the dimension of the eigenspace corresponding to 0).

The abbreviations "$\xrightarrow{a.s.}$" and "$\xrightarrow{i.d.}$" mean "converges almost surely to," and "converges in distribution to," respectively, while the abbreviation "i.i.d." means "independent identically distributed." We also use the abbreviation "*w.p.*1." to mean "with probability 1." For two sequences $\{x_k\}$ and $\{y_k\}$, we use the abbreviation $x_k = O(y_k)$ to denote that there exists $c > 0$ such that $\|x_k\| \leq c\|y_k\|$ for all k; and we use the abbreviation $x_k = \Theta(y_k)$ to denote that there exist $c_1, c_2 > 0$ such that $c_1\|y_k\| \leq \|x_k\| \leq c_2\|y_k\|$ for all k.

Throughout the paper, and *in the absence of an explicit statement to the contrary, we assume that A is singular*. This is done for convenience, since some of our analysis (e.g., the nullspace decomposition of the subsequent Prop. 1) makes no sense if A is nonsingular, and it would be awkward to modify so that it applies to both the singular and the nonsingular cases. However, our methods and analytical results have evident (and simpler) counterparts for the nonsingular case.

2. Iterative methods for singular systems. In this section, we review the convergence properties of the deterministic iteration (1) and its stochastic counterpart (2), as well as related classical algorithms and large-scale applications.

2.1. *Deterministic iterative methods.* Consider the deterministic iteration

$$x_{k+1} = x_k - \gamma G(Ax_k - b). \tag{3}$$

For a given triplet (A, b, G), with $b \in \mathbf{R}(A)$, we say that this iteration is *convergent* if there exists $\overline{\gamma} > 0$ such that for all $\gamma \in (0, \overline{\gamma})$ and all initial conditions $x_0 \in \Re^n$, the sequence $\{x_k\}$ produced by the iteration converges to a solution of $Ax = b$. The following condition, a restatement of conditions given in various contexts in the literature (e.g., [Kel65], [You72], [Tan74], [Dax90], [WB11]), is both necessary and sufficient for the iteration to be convergent (see [WB11]).

Assumption 1.

(a) Each eigenvalue of GA either has a positive real part or is equal to 0.

(b) The dimension of $\mathbf{N}(GA)$ is equal to the algebraic multiplicity of the eigenvalue 0 of GA.

(c) $\mathbf{N}(A) = \mathbf{N}(GA)$.

The following proposition, proved in [WB11], gives a decomposition of GA that will be useful in subsequent analysis for both the deterministic and simulation-based iterative methods.

PROPOSITION 1 (Nullspace Decomposition). *Let Assumption 1 hold. Then the matrix GA can be written as*

$$GA = [U\ V] \begin{bmatrix} 0 & N \\ 0 & H \end{bmatrix} [U\ V]',$$

where

U is a matrix whose columns form an orthonormal basis of $\mathbf{N}(A)$.
V is a matrix whose columns form an orthonormal basis of $\mathbf{N}(A)^{\perp}$.
$N = U'GAV$.
$H = V'GAV$, and its eigenvalues are equal to the eigenvalues of GA that have positive real parts.

The significance of the decomposition of Prop. 1 is that in a scaled coordinate system defined by the transformation

$$y = U'(x - x^*), \qquad z = V'(x - x^*),$$

where x^* is a solution of $Ax = b$, the iteration (3) decomposes into two component iterations, one for y, generating a sequence $\{y_k\}$, and one for z, generating a sequence $\{z_k\}$. The iteration decomposition can be written as

$$x_k = x^* + Uy_k + Vz_k,$$

where y_k and z_k are given by

$$y_k = U'(x_k - x^*), \qquad z_k = V'(x_k - x^*),$$

and are generated by the iterations

$$y_{k+1} = y_k - \gamma N z_k, \qquad z_{k+1} = z_k - \gamma H z_k. \tag{4}$$

Moreover the corresponding residuals $r_k = Ax_k - b$ are given by

$$r_k = AVz_k. \tag{5}$$

By analyzing the two iterations for y_k and z_k separately, the following result has been shown in [WB11].

PROPOSITION 2. *Assumption 1 holds if and only if iteration (3) is convergent, with $\gamma \in (0, \overline{\gamma})$ where*

$$\overline{\gamma} = \min \left\{ \frac{2a}{a^2 + b^2} \;\middle|\; a + bi \text{ is an eigenvalue of } GA, \quad a > 0 \right\}, \tag{6}$$

and the limit of iteration (3) *is the following solution of* $Ax = b$*:*

$$\hat{x} = (UU' - UNH^{-1}V')x_0 + (I + UNH^{-1}V')x^*, \tag{7}$$

where x_0 *is the initial iterate and* x^* *is the solution of* $Ax = b$ *that has minimum Euclidean norm.*

To outline the proof argument, let us write the component sequences $\{y_k\}$ and $\{z_k\}$ as

$$y_k = y_0 - \gamma N \sum_{t=0}^{k-1} (I - \gamma H)^t z_0, \qquad z_k = (I - \gamma H)^k z_0.$$

According to Prop. 1, the eigenvalues of H are equal to the eigenvalues of GA that have positive real parts. Let $a + bi$ be any such eigenvalue, and let γ be any scalar within the interval $(0, \overline{\gamma})$. By using Eq. (6) we have

$$0 < \gamma < \frac{2a}{a^2 + b^2},$$

or equivalently,

$$|1 - \gamma(a + bi)| < 1.$$

Therefore by taking $\gamma \in (0, \overline{\gamma})$, all eigenvalues of $I - \gamma H$ are strictly contained in the unit circle. In fact, we have $\rho(I - \gamma H) < 1$ if and only if $\gamma \in (0, \overline{\gamma})$, and $\rho(I - \gamma GA) \leq 1$ if and only if $\gamma \in (0, \overline{\gamma}]$. Thus for all $\gamma \in (0, \overline{\gamma})$ there exists an induced norm $\|\cdot\|_P$ such that $\|I - \gamma H\|_P < 1$. Therefore $z_k \to 0$, while y_k involves a convergent geometric series of powers of $I - \gamma H$, so it converges. The limit of x_k turns out to be the vector given by Eq. (7).

2.2. *Classical algorithms and applications.* We will now discuss some classical algorithms and problem types for which Assumption 1 is satisfied. Because this assumption is necessary and sufficient for the convergence of iteration (3) for some $\gamma > 0$, any set of conditions under which this convergence has been shown in the literature implies Assumption 1. In this section we collect various conditions of this type, which correspond to known algorithms of the form (3) or generalizations thereof. These fall in three categories: projection algorithms, proximal algorithms, and splitting algorithms.

Generally, these algorithms are used for finding a solution x^*, within a closed convex set X, of a variational inequality of the form

$$f(x^*)'(x - x^*) \geq 0, \qquad \forall\, x \in X, \tag{8}$$

mostly for cases where f is a mapping that is monotone on X, in the sense that $\big(f(y)-f(x)\big)'(y-x)\geq 0$, for all x, $y\in X$. For the special case where $f(x)=Ax-b$ and $X=\Re^n$, strong (or weak) monotonicity of f is equivalent to $A\succ 0$ (or $A\succeq 0$), and these algorithms take the form (3) for special choices of γ and G.

The projection algorithm is obtained when G is positive definite symmetric, and is related to Richardson's method (see e.g., Hageman and Young [HY81]). The convergence of the projection method for the variational inequality (8) generally requires strong monotonicity of f (see Sibony [M.70]; also textbook discussions in Bertsekas and Tsitsiklis [BN89], Section 3.5.3, or Facchinei and Pang [FP03], Section 12.1). When translated to the special case where $f(x)=Ax-b$ and $X=\Re^n$, the conditions for convergence are that $A\succ 0$, G is positive definite symmetric, and the stepsize γ is small enough. A variant of the projection method for solving weakly monotone variational inequalities is the extragradient method of Korpelevich [Kor76] [see case (ii) of Prop. 3],[1] which allows the use of a projection-like iteration when $A\succeq 0$ (rather than $A\succ 0$). A special case where f is weakly monotone [it has the form $f(x)=\Phi'\bar{f}(\Phi x)$ for some strongly monotone mapping $\bar{f}$] and the projection method is convergent was given by Bertsekas and Gafni [BG82] [see case (i) of Prop. 3, which is slightly more general].

The proximal algorithm, often referred to as the "proximal point algorithm," uses

$$G=(A+\beta I)^{-1},$$

with $\gamma\in(0,1]$ and $\beta>0$. An interesting special case arises when the algorithm is applied to the system $A'\Sigma^{-1}Ax=A'\Sigma^{-1}b$, with Σ positive semidefinite symmetric, which is equivalent to $Ax=b$ for A not necessarily positive semidefinite, as long as $Ax=b$ has a solution. Then we obtain the method $x_{k+1}=x_k-\gamma G(Ax_k-b)$, where $\gamma\in(0,1]$ and

$$G=(A'\Sigma^{-1}A+\beta I)^{-1}A'\Sigma^{-1}.$$

The proximal algorithm has been analyzed extensively by Rockafellar [Roc76] for the variational inequality (8) (and more general problems), and subse-

[1]The extragradient method for solving the system $Ax=b$ with $A\succeq 0$ is usually described as follows: at the current iterate x_k, it takes the intermediate step $\hat{x}_k=x_k-\beta(Ax_k-b)$ where β is a sufficiently small positive scalar, and then takes the step $x_{k+1}=x_k-\gamma(A\hat{x}_k-b)$, which can also be written as $x_{k+1}=x_k-\gamma(I-\beta A)(Ax_k-b)$. This corresponds to $\overline{G}=I$ in part (ii) of Prop. 3. For a convergence analysis, see the original paper [Kor76], or more recent presentations such as [BN89], Section 3.5, or [FP03], Section 12.1.2. The terms "projection" and "extragradient" are not very well-suited for our context, since our system $Ax=b$ is unconstrained and need not represent a gradient condition of any kind.

quently by several other authors. It is well-known (Martinet [Mar70], Rockafellar [Roc76]) that when f is weakly monotone, the proximal algorithm is convergent [see cases (iii)-(iv) of Prop. 3]. A recent work by Kannan and Shanbhag [KS13] proposes a variant of the proximal algorithm, the so called iterative proximal point method, which applies iterative projection steps towards the proximal problem with changing centers and uses diminishing stepsizes to guarantee convergence.

Splitting algorithms apply to many practical contexts of solving linear systems, and can be shown to converge under various assumptions. For example, if A is positive semidefinite symmetric, (B, C) is a regular splitting of A (i.e. $B + C = A$ and $B - C \succ 0$), and $G = B^{-1}$, the algorithm

$$x_{k+1} = x_k - B^{-1}(Ax_k - b),$$

converges to a solution, as shown by Luo and Tseng [LP89]. Convergent Jacobi and asynchronous or Gauss-Seidel iterations are also well known in dynamic programming, where they are referred to as *value iteration* methods (see e.g., [Ber12], [Put94]). In this context, the system to be solved has the form $x = g + Px$, with P being a substochastic matrix, and under various assumptions on P, the iteration

$$x_{k+1} = x_k - \gamma\big((I - P)x_k - g\big), \tag{9}$$

can be shown to converge asynchronously to a solution for some $\gamma \in (0, 1]$. Also asynchronous and Gauss-Seidel versions of iterations of the form (9) are known to be convergent, assuming that the matrix P has nonnegative entries and is irreducible, with $\rho(P) = 1$ (see [BN89], p. 517). In the special case where P or P' is an irreducible transition probability matrix and $g = 0$, the corresponding system, $x = Px$, contains as special cases the problems of consensus (multi-agent agreement) and of finding the invariant distribution of an irreducible Markov chain (see [BN89] Sections 7.3.1-7.3.2). Finally, a Gauss-Seidel algorithm is also known to be convergent for the case where P is irreducible and weakly diagonally dominant (see [BN89], Section 7.2.2). Under these assumptions, the matrix $A = I - P$, as well as the matrix G (equaling I, or resulting from A by using a Gauss-Seidel splitting), satisfy Assumption 1.

The following proposition collects various sets of conditions under which Assumption 1 holds. Some of these conditions can be shown by applying Prop. 2 in conjunction with known convergence results that have appeared in the literature just cited. In particular, parts (iii), (iv) and (v) of the following proposition can be shown in this way. We give independent proofs

which provide some special insights and set the stage for the simulation-based extensions to be addressed later. Parts (i), (ii) may not be shown by using known results because the convergence analysis of [BG82] and [Kor76] applies under somewhat more restrictive conditions. Of course Assumption 1 is more general than the union of the sets of conditions given in the proposition. For example it is satisfied if $GA \succ 0$, or more generally, if G is invertible, and GA is positive semidefinite with no purely imaginary eigenvalues. An example of a case where Assumption 1 is violated with GA invertible and $GA \succeq 0$ is when A is a 2×2 orthogonal rotation matrix and $G = I$ (GA has imaginary eigenvalues). It can then be verified that for $b = 0$, the sequence of iterates $x_k = (I - \gamma A)^k x_0$ diverges for any $\gamma > 0$ starting from any x_0 other than the unique solution $x^* = 0$; this is a classical example (see e.g., [BN89], Section 3.5, or [FP03], Example 12.1.3).

PROPOSITION 3. *Assumption 1 is satisfied if any one of the following conditions hold:*

(i) $A = \Phi' M \Phi$, *where* Φ *is an* $m \times n$ *matrix,* M *is an* $m \times m$ *matrix such that* $w' M w > 0$ *for all* $w \in \mathbf{R}(\Phi)$ *with* $w \neq 0$, *and* G *is a symmetric matrix such that* $v' G v > 0$ *for all* $v \in \mathbf{N}(\Phi)^\perp$ *with* $v \neq 0$.

(ii) $G = \overline{G} - \beta \overline{G} A \overline{G}$, *where* $\overline{G}$ *is an invertible matrix such that* $\overline{G} A \succeq 0$ *and* β *is a sufficiently small positive scalar.*

(iii) $A \succeq 0$ *and* $G = (A + \beta I)^{-1}$, *where* β *is any positive scalar.*

(iv) $G = (A' \Sigma^{-1} A + \beta I)^{-1} A' \Sigma^{-1}$, *where* $\Sigma \succ 0$ *is symmetric and* β *is any positive scalar.*

(v) $A \succeq 0$ *is symmetric,* (B, C) *is a regular splitting of* A *(i.e.* $B + C = A$ *and* $B - C \succ 0$*), and* $G = B^{-1}$.

PROOF. (i) First we claim that

$$\mathbf{N}(\Phi) = \mathbf{N}(A) = \mathbf{N}(A') = \mathbf{N}(GA). \tag{10}$$

Indeed, if x is such that $Ax = 0$ and $x \neq 0$, we have $x' A x = 0$ or equivalently $x' \Phi' M \Phi x = 0$. Since by assumption, $(\Phi x)' M (\Phi x) > 0$ for $\Phi x \neq 0$, we must have $\Phi x = 0$ and it follows that $\mathbf{N}(A) \subset \mathbf{N}(\Phi)$. Also if $\Phi x = 0$, then $Ax = 0$, showing that $\mathbf{N}(A) \supset \mathbf{N}(\Phi)$. Thus we have $\mathbf{N}(A) = \mathbf{N}(\Phi)$, and the same argument can be applied to $A' = \Phi' M' \Phi$ to show that $\mathbf{N}(A') = \mathbf{N}(\Phi) = \mathbf{N}(A)$. Finally, to show that $\mathbf{N}(GA) = \mathbf{N}(A)$, note that clearly $\mathbf{N}(GA) \supset \mathbf{N}(A)$. For the reverse inclusion, if $x \in \mathbf{N}(GA)$, we have $Gv = 0$, where $v = Ax$, so $v \in \mathbf{R}(A) = \mathbf{N}(A')^\perp = \mathbf{N}(\Phi)^\perp$. Thus by the assumption on G, we must have $v = 0$ and hence $x \in \mathbf{N}(A)$.

We will derive the decomposition of Prop. 1 by letting U and V be the orthonormal bases of $\mathbf{N}(A)$ and $\mathbf{N}(A)^\perp$ respectively. We have

(11)
$$[U\ V]'GA[U\ V] = \begin{bmatrix} U'GAU & U'GAV \\ V'GAU & V'GAV \end{bmatrix} = \begin{bmatrix} 0 & U'GAV \\ 0 & V'GAV \end{bmatrix} = \begin{bmatrix} 0 & N \\ 0 & H \end{bmatrix},$$

Consider the matrix $H = V'GAV$. We have

(12)
$$\begin{aligned} H = V'GAV &= V'G[U\ V][U\ V]'AV \\ &= V'G(UU' + VV')AV \\ &= (V'GV)(V'AV), \end{aligned}$$

where the last equality uses the fact $\mathbf{N}(A) = \mathbf{N}(A')$ shown earlier, which implies that $A'U = 0$ and hence $U'A = 0$. The assumption on G implies that the matrix $V'GV$ is symmetric positive definite, so it can be written as $V'GV = DD$ where D is symmetric positive definite. Thus from Eq. (12), $V'GAV$ is equal to $DD(V'AV)$, which is in turn similar to $D(V'AV)D$.

Since V is an orthonormal basis of $\mathbf{N}(GA)^\perp = \mathbf{N}(A)^\perp = \mathbf{N}(\Phi)^\perp$, the matrix ΦV has independent columns that belong to $\mathbf{R}(\Phi)$, so $V'AV = (\Phi V)'M(\Phi V) \succ 0$. It follows that $D(V'AV)D \succ 0$, so $D(V'AV)D$ has eigenvalues with positive real parts, and the same is true for the similar matrix $H = V'GAV$. Thus, from Eq. (11), GA has eigenvalues that either are equal to 0 or have positive real parts, and the algebraic multiplicity of the 0 eigenvalue of GA equals the dimension of $\mathbf{N}(GA)$.

(ii) We note that

$$GA = \overline{G}A - \beta(\overline{G}A)^2.$$

The idea of the proof is that the term $-\beta(\overline{G}A)^2$ adds a positive real part to any purely imaginary eigenvalues of $\overline{G}A$, thereby satisfying Assumption 1(a). Indeed, each eigenvalue λ of GA has the form $\lambda = \mu - \beta\mu^2$, where μ is a corresponding eigenvalue of $\overline{G}A$. Since $\overline{G}A \succeq 0$, either $\mu = 0$ in which case $\lambda = 0$, or μ has positive real part, in which case the same is true for λ where β is sufficiently small, or μ is purely imaginary in which case the real part of λ is $\beta|\mu|^2$ and is positive. Thus Assumption 1(a) holds for β sufficiently small. Also for β sufficiently small, the algebraic multiplicity of 0 eigenvalue of $\overline{G}A$ is equal to the algebraic multiplicity of 0 eigenvalue of GA. Since $\overline{G}A \succeq 0$, the algebraic multiplicity of 0 eigenvalue of $\overline{G}A$ is equal to the dimension of $\mathbf{N}(\overline{G}A)$,[2] which is less or equal to the dimension of $\mathbf{N}(GA)$. It follows that the algebraic multiplicity of a 0 eigenvalue of GA is less or equal

[2]For an arbitrary matrix $A \succeq 0$, the algebraic multiplicity of the eigenvalue 0 is equal to the dimension of $\mathbf{N}(A)$. For a proof, note that if this is not so, there exists v such that

to the dimension of $\mathbf{N}(GA)$, so it must be equal, and Assumption 1(b) holds. Finally, $G = \overline{G}(I - \beta A\overline{G})$, where $\overline{G}$ is invertible and $I - \beta A\overline{G}$ is invertible for sufficiently small β, so G is invertible and Assumption 1(c) holds.

(iii) Let $A = WQW'$ be a Schur decomposition of A, where W is unitary and Q is upper triangular with the eigenvalues of A along its diagonal (for the existence and properties of a Schur decomposition, see references on linear algebra, such as [Bel70], Chapter 11, [Gol91], Section 6.3, or [TB97], Lecture 24). We have

$$GA = (A + \beta I)^{-1}A = (WQW' + \beta I)^{-1}WQW' = W(Q + \beta I)^{-1}QW'.$$

Note that $(Q + \beta I)^{-1}$ is the inverse of an upper triangular matrix so it is upper triangular, and $(Q + \beta I)^{-1}Q$ is the product of upper triangular matrices, so it is also upper triangular. Thus we have obtained a Schur decomposition of GA. Let $a + bi$ be an eigenvalue of A, which is also a diagonal entry of Q. Then the corresponding eigenvalue of GA is

$$\frac{a + bi}{a + \beta + bi} = \frac{a^2 + a\beta + b^2 + \beta bi}{(a + \beta)^2 + b^2}.$$

Since $A \succeq 0$, each eigenvalue of A has nonnegative real part. It follows that each eigenvalue of GA either is equal to 0 ($a = b = 0$) or has positive real part ($a^2 + a\beta + b^2 > 0$ if $a \neq 0$ or $b \neq 0$). Thus Assumption 1(a) is satisfied. Moreover, since G is invertible, $\mathbf{N}(GA) = \mathbf{N}(A)$ [cf. Assumption 1(c)]. Also note that an eigenvalue of GA equals 0 if and only if the corresponding eigenvalue of A equals 0, which implies that the algebraic multiplicities of the eigenvalue 0 for GA and A are equal. It follows that Assumption 1(b) is satisfied since $A \succeq 0$, and the algebraic multiplicity of the 0 eigenvalue of any positive semidefinite matrix is equal to the dimension of its nullspace.

(iv) We have

$$GA = (A'\Sigma^{-1}A + \beta I)^{-1}A'\Sigma^{-1}A = (\overline{A} + \beta I)^{-1}\overline{A}, \tag{13}$$

where $\overline{A} = A'\Sigma^{-1}A$. Thus GA has the form of case (iii), and hence satisfies conditions (a) and (b) of Assumption 1. To prove condition (c), note that from Eq. (13) we have $\mathbf{N}(GA) = \mathbf{N}(A'\Sigma^{-1}A)$, while from the argument

$Av \neq 0$ and $A^2v = 0$. Let $u = Av$ so that $Au = A^2v = 0$. Now for any $\beta > 0$ we have

$$(u - \beta v)'A(u - \beta v) = -\beta u'Av + \beta^2 v'Av = -\beta u'u + \beta^2 v'u.$$

By taking β to be sufficiently close to 0 we have $(u - \beta v)'A(u - \beta v) < 0$, arriving at a contradiction.

of case (i) [cf. Eq. (10)], $\mathbf{N}(A'\Sigma^{-1}A) = \mathbf{N}(A)$. Thus, $\mathbf{N}(GA) = \mathbf{N}(A)$ and Assumption 1(c) is satisfied.

(v) Using the facts $A \succeq 0$ and $B - C \succ 0$, we have $2B = A + (B - C) \succ 0$ so that $B \succ 0$. By the symmetry of $A \succeq 0$ we have $A = A^{1/2}A^{1/2}$ where $A^{1/2} \succeq 0$ is also symmetric. Let the diagonal decomposition of $A^{1/2}$ be

$$A^{1/2} = [U\ V] \begin{bmatrix} 0 & 0 \\ 0 & \Lambda \end{bmatrix} [U\ V]'$$

where Λ is positive diagonal, V is an orthonormal basis of $\mathbf{N}(A)^{\perp} = \mathbf{R}(A)$, and U is an orthonormal basis of $\mathbf{N}(A)$. Let $E = UU'$ and we have

$$A^{1/2} + E = [U\ V] \begin{bmatrix} I & 0 \\ 0 & \Lambda \end{bmatrix} [U\ V]',$$

as well as

$$A^{1/2}(A^{1/2} + E)^{-1} = [U\ V] \begin{bmatrix} 0 \cdot I & 0 \\ 0 & \Lambda^{-1}\Lambda \end{bmatrix} [U\ V]' = VV'.$$

By using the invertibility of $A^{1/2} + E$, the fact $A^{1/2}VV' = A^{1/2}$, and the preceding relations, we write

$$\begin{aligned} GA = B^{-1}A &= B^{-1}A^{1/2}A^{1/2} \\ &\sim (A^{1/2} + E)B^{-1}A^{1/2}A^{1/2}(A^{1/2} + E)^{-1} \\ &= (A^{1/2} + E)B^{-1}A^{1/2}VV' \\ &= A^{1/2}B^{-1}A^{1/2} + EB^{-1}A^{1/2}. \end{aligned}$$

By using the diagonal decomposition of $A^{1/2}$ we further obtain

$$A^{1/2}B^{-1}A^{1/2} = [U\ V] \begin{bmatrix} 0 & 0 \\ 0 & \Lambda V'B^{-1}V\Lambda \end{bmatrix} [U\ V]',$$

Since $B \succ 0$ and $V\Lambda^{-1}$ is full rank, we have $\Lambda^{-1}V'BV\Lambda^{-1} \succ 0$ and its eigenvalues have positive real parts. It follows that the eigenvalues of its inverse $\Lambda V'B^{-1}V\Lambda$ also have positive real parts. Also note that $V'E = 0$ and $A^{1/2}U = 0$, so we have

$$\begin{aligned} EB^{-1}A^{1/2} &= [U\ V] \begin{bmatrix} U'EB^{-1}A^{1/2}U & U'EB^{-1}A^{1/2}V \\ V'EB^{-1}A^{1/2}U & V'EB^{-1}A^{1/2}V \end{bmatrix} [U\ V]' \\ &= [U\ V] \begin{bmatrix} 0 & L \\ 0 & 0 \end{bmatrix} [U\ V]', \end{aligned}$$

where $L = U'E^{-1}A^{1/2}V$. Finally, we have

$$\begin{aligned} GA &\sim [U\ V] \begin{bmatrix} 0 & 0 \\ 0 & \Lambda V' B^{-1} V \Lambda \end{bmatrix} [U\ V]' + [U\ V] \begin{bmatrix} 0 & L \\ 0 & 0 \end{bmatrix} [U\ V]' \\ &\sim \begin{bmatrix} 0 & L \\ 0 & \Lambda V' B^{-1} V \Lambda \end{bmatrix}. \end{aligned} \tag{14}$$

According to Eq. (14), eigenvalues of GA either have positive real parts or are equal to 0. Also the number of 0 eigenvalues of GA is equal to the dimension of $\mathbf{N}(A^{1/2}) = \mathbf{N}(A)$. Therefore GA satisfies parts (a) and (b) of Assumption 1. Finally since $G = B^{-1}$ is invertible, Assumption 1(c) is also satisfied. □

The following examples describe several interesting applications where Prop. 3 applies.

Example 1 (Overdetermined Least Squares Problem). Consider the weighted least squares problem

$$\min_{x \in \Re^n} \|Cx - d\|_\xi^2,$$

where C is an $m \times n$ matrix with $m > n$, and $\|\cdot\|_\xi$ is a weighted Euclidean norm with ξ being a vector with positive components, i.e. $\|y\|_\xi^2 = \sum_{i=1}^m \xi_i y_i^2$. This problem is equivalent to the $n \times n$ system $Ax = b$ where

$$A = C'\Xi C, \qquad b = C'\Xi d,$$

and Ξ is the diagonal matrix that has the components of ξ along the diagonal. Here A is symmetric positive semidefinite, and with the choices of G given by Prop. 3(i),(iii),(iv),(v), Assumption 1 is satisfied.

The following examples involve the approximation of the solution of a high-dimensional problem within a lower dimensional subspace

$$S = \{\Phi x \mid x \in \Re^n\},$$

where Φ is an $m \times n$ matrix whose columns can be viewed as features/basis functions, in the spirit of Galerkin and Petrov-Galerkin approximation (see e.g., Krasnoselskii et. al. [Kra72], Saad [Saa03]).

Example 2 (Least Squares with Subspace Approximation). Consider the least squares problem

$$\min_{y \in \Re^m} \|Cy - d\|_\xi^2,$$

where C and d are given $s\times m$ matrix and vector in $\Re^s$, respectively, and $\|\cdot\|_\xi$ is the weighted Euclidean norm of Example 1. By approximating y within the subspace $S = \{\Phi x \mid x \in \Re^n\}$, we obtain the least squares problem

$$\min_{x\in\Re^n} \|C\Phi x - d\|_\xi^2.$$

This is equivalent to the $n \times n$ linear system $Ax = b$ where

$$A = \Phi' C' \Xi C \Phi, \qquad b = \Phi' C' \Xi d.$$

Similar to Example 1, A is symmetric positive semidefinite. With the choices of G given by Prop. 3(i),(iii),(iv), (v), Assumption 1 is satisfied. Simulation-based noniterative methods using the formulation of this example were proposed in Bertsekas and Yu [BY09], while iterative methods were proposed in Wang et al. [WPB09] and tested on large-scale inverse problems in Polydorides et al. [PWB10]. The use of simulation may be desirable in cases where either s or m, or both, are much larger than n. In such cases the explicit calculation of A may be difficult.

The next two examples involve a nonsymmetric matrix A. They arise in the important context of policy iteration in approximate dynamic programming (ADP for short) as well as in general Galerkin approximation [Kra72]; see e.g., the books [BN96], [SB98], [Pow11] and [Ber12], and the survey [Ber11].

Example 3 (Projected Equations with Subspace Approximation). Consider a projected version of an $m \times m$ fixed point equation $y = Py + g$ given by

$$\Phi x = \Pi_\xi(P\Phi x + g),$$

where Π_ξ denotes orthogonal projection onto the subspace S with respect to the weighted Euclidean $\|\cdot\|_\xi$ of Examples 1 and 2 (in ADP, P is a substochastic matrix and g is the one-stage cost vector). By writing the orthogonality condition for the projection, it can be shown that this equation is equivalent to the $n \times n$ system $Ax = b$ where

$$A = \Phi' \Xi (I - P)\Phi, \qquad b = \Phi' \Xi g.$$

Various conditions guaranteeing that $A \succeq 0$ or $A \succ 0$ are given in [BY09] and [Ber11], and they involve contraction properties of the mappings $\Pi_\xi P$ and P. Examples are standard Markov and semi-Markov decision problems, where $y'\Xi(I - P)y > 0$ for all $y \in \mathbf{R}(\Phi)$ with $y \neq 0$ and an appropriate choice of Ξ, and $A \succeq 0$, so with the choices of G given by Prop. 3(i),(iii),(iv), Assumption 1 is satisfied.

EXAMPLE 4 (Oblique Projected Equations and Aggregation). The preceding example of projected equations $\Phi x = \Pi_\xi(g + Px)$ can be generalized to the case where Π_ξ is an oblique projection, i.e., its range is $S = \{\Phi x \mid x \in \Re^n\}$ and is such that $\Pi_\xi^2 = \Pi_\xi$. Let Ψ be an $m \times n$ matrix such that $\mathbf{R}(\Psi)$ does not contain any vector orthogonal to $\mathbf{R}(\Phi)$, and let Π_ξ be the weighted oblique projection such that $\Pi_\xi y \in S$ and $(y - \Pi_\xi y)'\Xi\Psi = 0$ for all $y \in \Re^m$. The optimality condition associated with the projected equation is

$$\Psi'\Xi\Phi x = \Psi'\Xi(g + P\Phi x),$$

which is equivalent to the $n \times n$ system $Ax = b$ where

$$A = \Psi'\Xi(I - P)\Phi, \qquad b = \Psi'\Xi g.$$

We don't necessarily have $A \succ 0$ or $A \succeq 0$ even if P is a substochastic matrix. With the choice of G given by Prop. 3(iv), Assumption 1 is satisfied.

One special case where oblique projections arise in ADP is an aggregation equation of the form

$$\Phi x = \Phi D(g + \alpha P\Phi x),$$

where $\alpha \in (0, 1]$, D is an $n \times m$ matrix, and the n-dimensional rows of Φ and the m-dimensional rows of D are probability distributions (see [Ber12] Section 7.3.6, and Yu and Bertsekas [YB12]). Assume that for a collection of n mutually disjoint subsets of the index set $\{1, \ldots, m\}$, $I_1, \ldots, I_m$, we have $d_{ji} > 0$ only if $i \in I_j$ and $\phi_{ij} = 1$ if $i \in I_j$. Then it can be verified that $D\Phi = I$ hence $(\Phi D)^2 = \Phi D$, so ΦD is an oblique projection matrix. The aggregation equation is equivalent to the $n \times n$ system $Ax = b$ where

$$A = I - \alpha DP\Phi, \qquad b = Dg.$$

In standard discounted problems, $DP\Phi$ is a stochastic matrix and $\alpha < 1$. Then iteration (9) where $G = I$ and $\gamma \in (0, 1]$ is convergent. For additional choices of G such that Assumption 1 is satisfied, we refer to the discussion following Eq. (9).

2.3. *Simulation-based methods.* We will now consider a simulation-based version of the deterministic method (3). It has the form

$$x_{k+1} = x_k - \gamma G_k(A_k x_k - b_k), \tag{15}$$

where A_k, b_k, and G_k, are estimates of A, b, and G, respectively. Throughout our analysis, we assume the following.

ASSUMPTION 2. *The sequence* $\{A_k, b_k, G_k\}$ *is generated by a stochastic process such that*

$$A_k \xrightarrow{a.s.} A, \qquad b_k \xrightarrow{a.s.} b, \qquad G_k \xrightarrow{a.s.} G.$$

Assumption 2 is a general assumption that applies to practical situations involving a stochastic simulator/sampler. In many of these applications, the simulation process generates an infinite sequence of random variables

$$\big\{(W_t, v_t) \mid t = 1, 2, \dots\big\},$$

where W_t is an $n \times n$ matrix and v_t is a vector in $\Re^n$, and estimates A and b with A_k and b_k given by

$$A_k = \frac{1}{k}\sum_{t=1}^{k} W_t, \qquad b_k = \frac{1}{k}\sum_{t=1}^{k} v_t. \tag{16}$$

For instance, the sample sequence may consist of independent samples from a certain distribution (e.g., Drineas et al. [DMM06]) or from a sequence of importance sampling distributions. Also, the sample sequence can be generated through state transitions of an irreducible Markov chain, as for example in temporal difference methods in the context of ADP (e.g., [BB96], [Boy02], [NB03], and [Ber10]), or for general projected equations (e.g., [BY09], [Ber11]).

Stochastic algorithms that use Monte Carlo estimates of the form (16) have a long history in stochastic programming and applies to a wide range of problems under various names (for recent theoretical developments, see Shapiro [Sha03], and for applications in ADP, see [BN96]). The proposed method in the current work uses increasingly accurate approximations, obtained from some sampling process, to replace unknown quantities in deterministic algorithms that are known to be convergent. A related method, known as the *sample average approximation method* (SAA), approximates the original problem by using a fixed number of samples obtained from presampling (e.g., see Shapiro et al. [SDR09] for a book account, and related papers such as Kleywegt et al. [KSHdM02], and Nemirovskii et al. [NJLS09]). A variant of SAA is the so-called retrospective approximation method (RA), which solves a sequence of SAA problems by using an increasing number of samples for each problem (see e.g., Pasupathy [Pas10]). Our method differs from RA in that our algorithm is an iterative one-time scale method that uses increasingly accurate approximations in the iteration, instead of solving a sequence of increasingly accurate approximate systems. Throughout this paper, we avoid defining explicitly $\{A_k, b_k, G_k\}$ as sample averages, so our analysis applies to a more general class of stochastic methods.

Another related major class of methods, known as the *stochastic approximation method* (SA), uses a single sample per iteration and a decreasing sequence of stepsizes $\{\gamma_k\}$ to ensure convergence (see e.g., [BMP90], [Bor08], [KY03], and [Mey07] for textbook discussions, and see Nemirovski et al. [NJLS09] for a recent comparison between SA and SAA). Our methodology differs in fundamental ways from SA. While the SA method relies on decreasing stepsizes to control the convergence process, our methodology is based on Monte-Carlo estimates and uses a constant stepsize, which implies a constant modulus of contraction as well as multiplicative (rather than additive) noise. This both enhances its performance and complicates its analysis when A is singular, as it gives rise to large stochastic perturbations that must be effectively controlled to guarantee convergence.

We will first illustrate some possibilities for obtaining $\{A_k, b_k, G_k\}$ by simulation, based on the applications of Examples 1-4. As noted earlier, the use of simulation in these applications aims to deal with large-scale linear algebra operations, which would be very time consuming or impossible if done exactly. In the first application we aim to solve approximately an overdetermined system by randomly selecting a subset of the constraints; see [DMM06], [DMMS11].

EXAMPLE 5 (Continued from Example 1). Consider the least squares problem of Example 1, which is equivalent to the $n \times n$ system $Ax = b$ where

$$A = C'\Xi C, \qquad b = C'\Xi d.$$

We generate a sequence of i.i.d. indices $\{i_1, \ldots, i_k\}$ according to a distribution ζ, and estimate A and b using Eq. (16), where

$$W_t = \frac{\xi_{i_t}}{\zeta_{i_t}} c_{i_t} c'_{i_t}, \qquad v_t = \frac{\xi_{i_t}}{\zeta_{i_t}} c_{i_t} d_{i_t},$$

c'_i is the ith row of C, ξ_i is the ith diagonal component of Ξ, and the probability ζ_i satisfies $\zeta_i > 0$ if $\xi_i > 0$.

EXAMPLE 6 (Continued from Example 2). Consider the least squares problem of Example 2, which is equivalent to the $n \times n$ system $Ax = b$ where

$$A = \Phi' C' \Xi C \Phi, \qquad b = \Phi' C' \Xi d.$$

We generate i.i.d. indices $\{i_1, \ldots, i_k\}$ according to a distribution ζ, and then generate two sequences of independent state transitions $\{(i_1, j_1), \ldots (i_k, j_k)\}$ and $\{(i_1, \ell_1), \ldots (i_k, \ell_k)\}$ according to transition probabilities p_{ij} (i.e., given

i_k, generate (i_k, j_k) with probability $p_{i_k j_k}$). We may then estimate A and b using Eq. (16), where

$$W_t = \frac{\xi_{i_t} c_{i_t j_t} c_{i_t \ell_t}}{\zeta_{i_t} p_{i_t j_t} p_{i_t \ell_t}} \phi_{j_t} \phi'_{\ell_t}, \qquad v_t = \frac{\xi_{i_t} c_{i_t j_t}}{\zeta_{i_t} p_{i_t j_t}} \phi_{i_t} d_{j_t},$$

ϕ'_i is the ith row of Φ, and c_{ij} is the (i,j)th component of C. The sampling probabilities need to satisfy that $p_{ij} > 0$ if $c_{ij} > 0$.

EXAMPLE 7 (Continued from Example 3). Consider the projected equation of Example 3, which is equivalent to the $n \times n$ system $Ax = b$ where

$$A = \Phi' \Xi (I - P) \Phi, \qquad b = \Phi' \Xi g.$$

One approach is to generate a sequence of i.i.d. indices $\{i_1, \ldots, i_k\}$ according to distribution ζ, and generate a sequence of state transitions $\{(i_1, j_1), \ldots (i_k, j_k)\}$ according to transition probabilities θ_{ij}. We may then estimate A and b using Eq. (16), where

$$W_t = \frac{\xi_{i_t}}{\zeta_{i_t}} \phi_{i_t} \left(\phi_{i_t} - \frac{p_{i_t j_t}}{\theta_{i_t j_t}} \phi_{j_t} \right)', \qquad v_t = \frac{\xi_{i_t}}{\zeta_{i_t}} \phi_{i_t} g_{i_t},$$

p_{ij} denotes the (i,j)th component of the matrix P. The sampling probabilities need to satisfy $\zeta_i > 0$ if $\xi_i > 0$ and $\theta_{ij} > 0$ if $p_{ij} > 0$.

In an alternative approach, which applies to cost evaluation of discounted ADP problems, the matrix P is the transition probability matrix of an irreducible Markov chain. We use the Markov chain instead of i.i.d. indices for sampling. In particular, we take ξ to be the invariant distribution of the Markov chain. We then generate a sequence $\{i_1, \ldots, i_k\}$ according to this Markov chain, and estimate A and b using Eq. (16), where

$$W_t = \phi_{i_t} (\phi_{i_t} - \phi_{i_{t+1}})', \qquad v_t = \phi_{i_t} g_{i_t}.$$

It can be verified that $A_k = \frac{1}{k} \sum_{t=1}^{k} W_t \xrightarrow{a.s.} A$ and $b_k = \frac{1}{k} \sum_{t=1}^{k} v_t \xrightarrow{a.s.} b$ by the strong law of large numbers for irreducible Markov chains.

EXAMPLE 8 (Continued from Example 4). Consider the projected equation using oblique projection of Example 4, which is equivalent to the $n \times n$ system $Ax = b$ where

$$A = \Psi' \Xi (I - P) \Phi, \qquad b = \Psi' \Xi g.$$

We may generate a sequence of i.i.d. indices $\{i_1, \ldots, i_k\}$ according to distribution ζ, generate a sequence of state transitions $\{(i_1, j_1), \ldots, (i_k, j_k)\}$

according to transition probabilities θ_{ij}, and estimate A and b using Eq. (16), where

$$W_t = \frac{\xi_{i_t}}{\zeta_{i_t}} \psi_{i_t} \left(\phi_{i_t} - \frac{p_{i_t j_t}}{\theta_{i_t j_t}} \phi_{j_t} \right)', \qquad v_t = \frac{\xi_{i_t}}{\zeta_{i_t}} \psi_{i_t} g_{i_t},$$

ψ_i' denotes the ith row of the matrix Ψ. The sampling probabilities need to satisfy $\zeta_i > 0$ if $\xi_i > 0$ and $\theta_{ij} > 0$ if $p_{ij} > 0$.

In the special case of the aggregation equation $\Phi x = \Phi D(g + \alpha P \Phi x)$ where P is a transition probability matrix, this is equivalent to $Ax = b$ where

$$A = I - \alpha D P \Phi, \qquad b = Dg.$$

We may generate i.i.d. indices $\{i_1, \ldots, i_k\}$ according to a distribution ζ, generate a sequence of state transitions $\{(i_1, j_1), \ldots, (i_k, j_k)\}$ according to P, and estimate A and b using Eq. (16), where

$$W_t = I - \frac{\alpha}{\zeta_{i_t}} d_{i_t} \phi_{j_t}', \qquad v_t = \frac{g_{i_t}}{\zeta_{i_t}} d_{i_t},$$

and d_i is the ith column of D. The sampling probabilities need to satisfy $\zeta_i > 0$ if $g_i > 0$.

Note that the simulation formulas used in Examples 5-8 satisfy Assumption 2, and only involve low-dimensional linear algebra computations. In Example 5, this is a consequence of the low dimension n of the solution space of the overdetermined system. In Examples 6-8, this is a consequence of the low dimension n of the approximation subspace defined by the basis matrix Φ.

Even if Assumptions 1, 2 are both satisfied, iteration (15) does not necessarily converge to any solution. To understand the reason, let us consider the decomposition of the iteration into the components Uy_k and Vz_k within $\mathbf{N}(A)$ and $\mathbf{N}(A)^\perp$, respectively (cf. Prop. 1). Then contrary to the case where there is no simulation error [cf. Eq. (4)], z_k is no longer decoupled from y_k, and may become contaminated by additional simulation noise through the y_k iterates, which are not governed by a contractive process. The decomposed iteration takes the form

$$y_{k+1} = y_k - \gamma N z_k + \zeta_k(y_k, z_k), \qquad z_{k+1} = z_k - \gamma H z_k + \xi_k(y_k, z_k), \tag{17}$$

where $\zeta_k(y_k, z_k)$ and $\xi_k(y_k, z_k)$ are simulation-induced errors that are functions of y_k and z_k [compare with Eq. (4)]. Generally, these errors converge

to 0 if $\{y_k\}$ and $\{z_k\}$ are bounded, in which case z_k converges to 0 (since $I-\gamma H$ is a contraction for appropriate γ by Prop. 1). However, y_k need not stay bounded, and as a result the sequence $\{x_k\}$ generally does not converge. On the other hand, the residual sequence $\{Ax_k - b\}$ is better behaved than $\{x_k\}$. An important fact in this regard is $Ax_k - b = A(x_k - x^*)$ (where x^* is any solution of $Ax = b$), so it depends only on the component of $x_k - x^*$ that belongs to $\mathbf{N}(A)^\perp$.

To address the divergence of the iteration, a class of modification/stabilization schemes for iteration (15) has been proposed in the related work [WB11], which aims to attenuate the effect of accumulating simulation errors. One such modification is

$$x_{k+1} = (1-\delta_k)x_k - \gamma G_k(A_k x_k - b_k), \tag{18}$$

whereby the eigenvalues of the iteration are shifted by $-\delta_k$, and $\delta_k \downarrow 0$ at a rate that is slower than the rate of convergence of $(G_kA_k - GA, b_k - b)$. It was shown that in this case, $\{x_k\}$ converges to a specific solution of $Ax = b$ with probability 1. This stabilization approach can be applied to an arbitrary stochastic iteration of the form (15), provided that the corresponding deterministic iteration is convergent. Another modification is the selective eigenvalue shifting scheme given by

$$x_{k+1} = \big(1 - \delta_k(I - \Pi_k)\big)x_k - \gamma G_k(A_k x_k - b_k), \tag{19}$$

where Π_k converges to the orthogonal projection matrix onto $\mathbf{N}(A)^\perp$, and $I - \Pi_k$ converges to the projection matrix onto $\mathbf{N}(A)$. This scheme requires an estimated projection matrix, and selectively shifts only those eigenvalues corresponding to $\mathbf{N}(A)$. It was proved that under reasonable assumptions, iteration (19) converges to the solution of $Ax = b$ with minimal Euclidean norm (the Moore-Penrose pseudoinverse solution). We will return to the problem of estimating Π_k in Section 6. A related work by Koshal et al. [KNS13] considers an iterative Tikhonov regularization method for stochastic variational inequalities, which uses a diminishing regularization term that has an effect similar to our stabilization term. This method differs from the method of Eqs. (18)-(19) in that it requires a diminishing stepsize γ_k, and it does not involve the multiplicative noise (i.e., A_k multiplying x_k in $A_kx_k - b_k$).

In the following sections, we will focus on the convergence of two algorithms of the form (15), for which it is unnecessary to use stabilization:

(a) When the nullspace of the iteration "remains stable," i.e.,

$$\mathbf{N}(G_kA_k) = \mathbf{N}(A_k) = \mathbf{N}(A) = \mathbf{N}(GA).$$

In this case, z_k does not depend on y_k and converges to 0, and the same is true for the residual $Ax_k - b = AVz_k$ [cf. Eqs. (4) and (5)]. Moreover, under some additional special conditions, we can show that y_k also converges. This analysis will be given in Section 3.

(b) When a proximal iteration involving quadratic regularization is applied to the system $A'\Sigma^{-1}Ax = A'\Sigma^{-1}b$. This iteration is given by

$$x_{k+1} = x_k - \left(A_k'\Sigma^{-1}A_k + \beta I\right)^{-1} A_k'\Sigma^{-1}(A_k x_k - b_k),$$

where β is a positive scalar. Here the special structure of the matrices $G_k = (A_k'\Sigma^{-1}A_k + \beta I)^{-1}A_k'\Sigma^{-1}$ is such that $I - G_kA_k$ is contractive for all k. Under an additional assumption on the convergence rate of $\{A_k, b_k, G_k\}$, this structure forces the residual to converge to 0. This analysis will be given in Section 4.

While in both of the above cases the sequence of residuals $\{Ax_k - b\}$ is convergent to 0 with probability 1, the sequence of iterates $\{x_k\}$ may be unbounded. In order to extract from $\{x_k\}$ a convergent sequence, we will propose in Section 6 two approaches for estimating the matrix of projection onto $\mathbf{N}(A)^\perp$, and we will apply them to the preceding algorithms.

2.4. *Choice of the stepsize γ.* An important issue related to our methodology is the selection of an appropriate stepsize γ in the simulation-based iterations to guarantee convergence. In theory, we have shown that any stepsize γ that works for the deterministic counterpart of the algorithm, will also work for the stochastic algorithm. More specifically, the stepsize γ needs to be sufficiently small, proportional to the smallest positive part of eigenvalues of GA (see Prop. 2). In practice and in the presence of simulation noise, determining γ can be challenging, given that A and b are largely unknown. This is particularly so for singular or nearly singular problems, since then the close-to-zero eigenvalues of GA are hard to estimate precisely. In this section, we will address this issue in several different cases.

There are a few special cases where we may infer an appropriate value of γ by using the general properties of G and A. As noted in Section 2.2, when the algorithm is a proximal point iteration, i.e.,

$$G_k = (A_k + \beta I)^{-1}, \qquad \text{or} \qquad G_k = (A_k'\Sigma^{-1}A_k + \beta I)^{-1}A_k'\Sigma^{-1},$$

where Σ is a positive definite symmetric matrix and β is a positive scalar, or is a splitting algorithm, we may simply take

$$\gamma = 1.$$

This stepsize is known to guarantee the convergence of the deterministic proximal and splitting iteration [e.g., cases (iii)-(v) in Prop. 3], and thus guarantee the convergence of the simulation-based iterations.

In more general cases, we may not have any access to G and A, and thus cannot choose γ easily as in the preceding cases. For these cases, one possibility is to replace the constant stepsize γ with a diminishing stepsize

$$\gamma_k \downarrow 0.$$

As long as $\sum_{k=0}^{\infty} \gamma_k = \infty$, our convergence analysis can be adapted to work with such stepsizes. This approach is simple, but it may be less desirable because it degrades the linear rate of convergence of the residuals of the associated algorithms, which is guaranteed if the stepsize is not diminished to 0.

To remedy this, we propose another possibility, which is to estimate a proper value of γ to satisfy $\gamma \in (0, \overline{\gamma})$, where $\overline{\gamma}$ is given by Eq. (6), based on the sampling process. This is a more appealing approach in that it preserves the rate of convergence of the algorithms, which we describe as follows.

To estimate an appropriate value of γ, we may update an *upper bound of stepsize* according to

$$\overline{\gamma}_{k+1} = \begin{cases} \overline{\gamma}_k & \text{if } \rho(I - \overline{\gamma}_k G_k A_k) \leq 1 + \delta_k, \\ \eta_k \overline{\gamma}_k & \text{if } \rho(I - \overline{\gamma}_k G_k A_k) > 1 + \delta_k, \end{cases} \tag{20}$$

where $\{\delta_k\}$ is a slowly diminishing positive sequence and $\{\eta_k\}$ is a sequence in $(0, 1)$, and choose the stepsize according to

$$\gamma_k \in (0, \overline{\gamma}_k).$$

Under suitable conditions, which guarantee that δ_k eventually becomes an upper bound of the maximum perturbation in eigenvalues of $G_k A_k$, we can verify that $\overline{\gamma}_k$ converges to some point within the interval $(0, \overline{\gamma}]$, as shown in the following proposition.

PROPOSITION 4. *Let Assumptions 1-2 hold, let $\{\eta_k\}$ be a sequence of positive scalars such that $\prod_{k=0}^{\infty} \eta_k = 0$, and let $\{\delta_k\}$ be a sequence of positive scalars such that $\delta_k \downarrow 0$ and $\epsilon_k/\delta_k \xrightarrow{a.s.} 0$, where*

$$\epsilon_k = \max_{i=1,\ldots,n} \big|\lambda_i(G_k A_k) - \lambda_i(GA)\big|,$$

with $\lambda_i(M)$ denoting the ith eigenvalue of a matrix M. Then, with probability 1, the sequence $\{\overline{\gamma}_k\}$ generated by iteration (20) converges to some value in $(0, \overline{\gamma}]$ within a finite number of iterations.

PROOF. By its definition, $\{\overline{\gamma}_k\}$ either converges to 0 or stays constant at some positive value for all k sufficiently large. Assume to arrive at a contradiction that $\overline{\gamma}_k$ eventually stays constant at some $\hat{\gamma} > \overline{\gamma}$, such that $\rho(I - \hat{\gamma}GA) > 1$ (cf. the analysis of Prop. 2). Note that for any $\gamma, k > 0$, we have

$$|\rho(I - \gamma G_k A_k) - \rho(I - \gamma GA)| \leq \gamma \max_{i=1,\ldots,n} \left|\lambda_i(G_k A_k) - \lambda_i(GA)\right| = \gamma \epsilon_k.$$

From the preceding relation, we have

$$\rho(I - \overline{\gamma}_k G_k A_k) = \rho(I - \hat{\gamma} G_k A_k) \geq \rho(I - \hat{\gamma} GA) - \hat{\gamma}\epsilon_k > 1 + \delta_k,$$

for sufficiently large k with probability 1, where we used the facts $\delta_k \downarrow 0$, $G_k A_k \xrightarrow{a.s.} GA$ (cf. Assumption 2), so that $\epsilon_k \xrightarrow{a.s.} 0$ (see e.g., the book on matrix perturbation theory by Stewart and Sun [SS90]). Thus according to iteration (20), $\overline{\gamma}_k$ needs to be decreased again, yielding a contradiction. It follows that $\overline{\gamma}_k$ eventually enters the interval $(0, \overline{\gamma}]$ such that $\rho(I - \gamma GA) \leq 1$ for all $\gamma \in (0, \overline{\gamma}]$, with probability 1.

Once $\overline{\gamma}_k$ enters the interval $(0, \overline{\gamma}]$, we have

$$\rho(I - \overline{\gamma}_k G_k A_k) \leq \rho(I - \overline{\gamma}_k GA) + \overline{\gamma}_k \epsilon_k \leq 1 + \delta_k,$$

for all k sufficiently large with probability 1, where we used the fact $\epsilon_k/\delta_k \xrightarrow{a.s.} 0$ and the boundedness of $\{\overline{\gamma}_k\}$. This together with Eq. (20) imply that $\overline{\gamma}_k$ eventually stays constant at a value within $(0, \overline{\gamma}]$. □

In the preceding approach, the error tolerance sequence $\{\delta_k\}$ needs to decrease more slowly than ϵ_k, the simulation errors in eigenvalues. Based on matrix perturbation theory, as $G_k A_k \xrightarrow{a.s.} GA$, we have $\epsilon_k \leq O(\|G_k A_k - GA\|^q)$, where $q = 1$ if GA is diagonalizable and $q = 1/n$ otherwise (see [SS90] p. 192 and p. 168). This allows us to choose δ_k in accordance with the convergence rate of the simulation error. In particular, when A_k and G_k are sample averages of random samples, these estimates often conform to a form of central limit theorem and certain concentration inequalities. For these cases, the eigenvalue perturbation ϵ_k decreases to zero at a rate of $O\left(k^{-q/2}\right)$ in some stochastic sense, so we can choose $\delta_k \gg k^{-q/2}$ to guarantee that $\epsilon_k/\delta_k \xrightarrow{a.s.} 0$. A typical situation is when the samples are i.i.d. Gaussian random variables, in which case we can choose the error tolerance sequence

$$\delta_k = k^{-q/2+\nu},$$

where ν is an arbitrary scalar in $(0, q/2)$, to satisfy the assumption of Prop. 4 and ensure the convergence of the estimated stepsize.[3]

Moreover, the sequence $\{\eta_k\}$ in Eq. (20) can be selected as, e.g.,

$$\eta_k = \eta \in (0,1), \qquad \text{or} \qquad \eta_k = 1 - 1/k,$$

as long as it ensures that γ_k can be decreased to arbitrarily small. Finally, accordingly to the preceding analysis, the stepsize γ_k can be selected to converge finitely to an appropriate stepsize value γ within $(0, \overline{\gamma})$ for k sufficiently large with probability 1. Thus our convergence analysis for constant stepsizes applies.

Note that computing the spectral radius $\rho(I - \bar{\gamma}_k G_k A_k)$ in Eq. (20) may not be prohibitive, especially for problems of moderate dimension that involve time-consuming simulation. Moreover, it is sufficient to update $\bar{\gamma}_k$ periodically instead of once per iteration, to economize the computation overhead.

The details of the extensions of our convergence analysis to the stepsize schemes described above are relatively simple and will not be given. To sum up, if γ cannot be properly chosen based on general properties of the corresponding deterministic algorithm, it can be estimated based on the sampling process to ensure convergence, or simply taken to be diminishing. In what follows, we will assume that γ is chosen to be a constant.

3. Nullspace-consistent simulation-based iterations. In this section, we consider a special case of the iterative method (15) under an assumption that parallels Assumption 1. It requires that the rank and the nullspace of the matrix $G_k A_k$ do not change as we pass to the limit. As a

[3] Consider the case where G_k and A_k are samples averages of i.i.d. Gaussian random variables, and $\delta_k = k^{-q/2+\nu} \gg k^{-q/2}$ for some $\nu \in (0, q/2)$. In this case, the random variable $G_k A_k$ can be viewed as a differentiable function of the sample averages. It can be seen that the sequence $\{\mathbf{E}\left[\|G_k A_k - GA\|^{2p} k^p\right]\}$ is bounded for any $p > 0$ (this is similar to Assumption 5, and can be proved using the argument of footnote 4.) Then we obtain that the sequence $\left\{\mathbf{E}[\epsilon_k^{2p/q} k^p]\right\}$ is bounded. Let ε be an arbitrary positive scalar. We use the Markov inequality to obtain

$$\sum_{k=0}^{\infty} \mathbf{P}\left(\frac{\epsilon_k}{\delta_k} > \varepsilon\right) \leq \sum_{k=0}^{\infty} \frac{1}{(\delta_k \varepsilon)^{2p/q}} \mathbf{E}\left[\epsilon_k^{2p/q}\right] \leq \varepsilon^{-2p/q} \sum_{k=0}^{\infty} (k \delta_k^{2/q})^{-p} = \sum_{k=0}^{\infty} O\left(k^{-2p\nu/q}\right) < \infty,$$

where the last inequality holds when we take p sufficiently large. It follows from the Borel-Cantelli lemma that $\epsilon_k/\delta_k < \varepsilon$ eventually with probability 1. Since ε is arbitrary, we have $\epsilon_k/\delta_k \xrightarrow{a.s.} 0$. Therefore the choice of $\delta_k = k^{-q/2+\nu}$ satisfies the assumption of Prop. 4 and works for the update rule (20) in the case of i.i.d. Gaussian samples. This analysis can be extended to general sampling processes with subgaussian distributions.

result the nullspace decomposition that is associated with GA (cf. Prop. 1) does not change as the iterations proceed.

ASSUMPTION 3.

(a) Each eigenvalue of GA either has a positive real part or is equal to 0.

(b) The dimension of $\mathbf{N}(GA)$ is equal to the algebraic multiplicity of the eigenvalue 0 of GA.

(c) With probability 1, there exists an index $\overline{k}$ such that

$$\mathbf{N}(A) = \mathbf{N}(A_k) = \mathbf{N}(G_k A_k) = \mathbf{N}(GA), \qquad \forall\ k \geq \overline{k}. \tag{21}$$

If the stochastic iteration (15) satisfies Assumption 3, we say that it is *nullspace-consistent*. Since Assumption 3 implies Assumption 1, the corresponding deterministic iteration $x_{k+1} = x_k - \gamma G(Ax_k - b)$ is convergent. The key part of the assumption, which is responsible for its special character in a stochastic setting, is part (c).

Let us describe an important special case where Assumption 3 holds. Suppose that A and A_k are of the form

$$A = \Phi' M \Phi, \qquad A_k = \Phi' M_k \Phi,$$

where Φ is an $m \times n$ matrix, M is an $m \times m$ matrix with $y'My > 0$ for all $y \in \mathbf{R}(\Phi)$ with $y \neq 0$, and M_k is a sequence of matrices such that $M_k \to M$. Examples 2, 3, 6, and 7 satisfy this condition. Assuming that G is invertible, we can verify that Assumption 3(c) holds. Moreover if G is positive definite symmetric, by using Prop. 3(i) we obtain that Assumption 3(a),(b) also hold. We will return to this special case in Section 3.2.

3.1. *Convergence of residuals.* We will now show that for nullspace-consistent iterations, the residual sequence $\{Ax_k - b\}$ always converges to 0, regardless of whether the iterate sequence $\{x_k\}$ converges or not. The idea is that, under Assumption 3, the matrices U and V of the nullspace decomposition of $I - \gamma GA$ remain unchanged as we pass to the nullspace decomposition of $I - \gamma G_k A_k$. This induces a favorable structure of the z_k-portion of the iteration, and decouples it from y_k [cf. Eq. (17)].

PROPOSITION 5 (Convergence of Residual for Nullspace-Consistent Iteration). *Let Assumptions 2 and 3 hold. Then there exists a scalar $\bar{\gamma} > 0$, such that for all $\gamma \in (0, \bar{\gamma}]$ and every initial iterate x_0, the sequence $\{x_k\}$ generated by iteration (15) satisfies $Ax_k - b \to 0$ and $A_k x_k - b_k \to 0$ with probability 1.*

PROOF. Let x^* be the solution of $Ax = b$ with minimal Euclidean norm. Then iteration (15) can be written as

$$x_{k+1} - x^* = (I - \gamma G_k A_k)(x_k - x^*) + \gamma G_k(b_k - A_k x^*). \tag{22}$$

In view of nullspace-consistency, the nullspace decomposition of GA of Prop. 1 can also be applied for nullspace decomposition of $I - \gamma G_k A_k$. Thus, with probability 1 and for sufficiently large k, we have

$$[U\ V]'\,(I - \gamma G_k A_k)\,[U\ V] = \begin{bmatrix} I & -\gamma U' G_k A_k V \\ 0 & I - \gamma V' G_k A_k V \end{bmatrix},$$

where the zero block in the second row results from the assumption $\mathbf{N}(A_k) = \mathbf{N}(A)$ [cf. Eq. (21)], so

$$V'(I - \gamma G_k A_k)U = V'U - \gamma V' G_k A_k U = V'U - 0 = 0.$$

Recalling the iteration decomposition

$$x_k = x^* + U y_k + V z_k,$$

we may rewrite iteration (22) as

$$\begin{bmatrix} y_{k+1} \\ z_{k+1} \end{bmatrix} = \begin{bmatrix} I & -\gamma U' G_k A_k V \\ 0 & I - \gamma V' G_k A_k V \end{bmatrix} \begin{bmatrix} y_k \\ z_k \end{bmatrix} + \begin{bmatrix} \gamma U' G_k e_k \\ \gamma V' G_k e_k \end{bmatrix}, \tag{23}$$

where $e_k = b_k - A_k x^*$. Note that the z_k-portion of this iteration is independent of y_k. Focusing on the asymptotic behavior of iteration (23), we observe that:

(a) The matrix $I - \gamma V' G_k A_k V$ converges almost surely to $I - \gamma V' G A V = I - \gamma H$, which is contractive for sufficiently small $\gamma > 0$ (cf. the proof of Prop. 2).

(b) $\gamma V' G_k e_k \xrightarrow{a.s.} 0$, because $G_k \xrightarrow{a.s.} G$ and $e_k \xrightarrow{a.s.} 0$.

Therefore the z_k-portion of iteration (23) is strictly contractive for k sufficiently large with additive error decreasing to 0, implying that $z_k \xrightarrow{a.s.} 0$. Finally, since $Ax_k - b = AVz_k$, it follows that $Ax_k - b \xrightarrow{a.s.} 0$. Moreover, we have

$$\begin{aligned} A_k x_k - b_k &= A_k(x_k - x^*) + (A_k x^* - b_k) \\ &= A_k(U y_k + V z_k) + (A_k x^* - b_k) \\ &= A_k V z_k + (A_k x^* - b_k) \\ &\xrightarrow{a.s.} 0, \end{aligned}$$

where the last equality uses the fact $A_k U = 0$. Thus we also have $A_k x_k - b_k \xrightarrow{a.s.} 0$. □

The proof of the preceding proposition shows that for a given stepsize $\gamma > 0$, the residual of the nullspace-consistent stochastic iteration

$$x_{k+1} = x_k - \gamma G_k(A_k x_k - b_k)$$

converges to 0 if and only if the matrix $I - \gamma GA$ is a contraction in $\mathbf{N}(A)^\perp$. This is also the necessary and sufficient condition for the residual sequence generated by the deterministic iteration

$$x_{k+1} = x_k - \gamma G(Ax_k - b)$$

to converge to 0 (and also for $\{x_k\}$ to converge to some solution of $Ax = b$).

Note that the sequence $\{x_k\}$ may diverge; see Example 12 in Section 5. To construct a convergent sequence, we note that by Assumption 3, $\mathbf{N}(A_k) = \mathbf{N}(A)$, so we can obtain the projection matrix from

$$\Pi_{\mathbf{N}(A)^\perp} = A_k'(A_k A_k')^\dagger A_k,$$

where $(A_k A_k')^\dagger$ is the Moore-Penrose pseudoinverse of $A_k A_k'$. Applying this projection to x_k yields the vector

$$\hat{x}_k = \Pi_{\mathbf{N}(A)^\perp} x_k = \Pi_{\mathbf{N}(A)^\perp}(x^* + Uy_k + Vz_k) = x^* + Vz_k,$$

where x^* is the minimum norm solution of $Ax = b$. Since $z_k \xrightarrow{a.s.} 0$, we have $\hat{x}_k \xrightarrow{a.s.} x^*$.

3.2. *Convergence of iterates.* We now turn to deriving conditions under which $\{x_k\}$ converges naturally. This requires that the first row in Eq. (23) has the appropriate asymptotic behavior.

Proposition 6 (Convergence of Nullspace-Consistent Iteration). *Let Assumptions 2 and 3 hold, and assume in addition that*

$$\mathbf{R}(G_k A_k) \subset \mathbf{N}(A)^\perp, \qquad G_k b_k \in \mathbf{N}(A)^\perp, \tag{24}$$

for k sufficiently large. Then there exists $\overline{\gamma} > 0$ such that for all $\gamma \in (0, \overline{\gamma}]$ and all initial iterates x_0, the sequence $\{x_k\}$ generated by iteration (15) converges to a solution of $Ax = b$ with probability 1.

Proof. With the additional conditions (24), we have

$$\gamma U' G_k A_k V = 0, \qquad \gamma U' G_k(b_k - A_k x^*) = 0,$$

for all k sufficiently large, so that the first row of Eq. (23) becomes $y_{k+1} = y_k$. Since Prop. 5 implies that $z_k \xrightarrow{a.s.} 0$, it follows that x_k converges with probability 1, and its limit is a solution of $Ax = b$. □

We now revisit the special case discussed following Assumption 3, and prove the convergence of $\{x_k\}$. This case arises in the context of ADP (cf. Examples 2, 3, 6, 7), and the convergence of iteration (15) within that context has been discussed in [Ber11]. It involves the approximation of the solution of a high-dimensional equation within a lower-dimensional subspace spanned by a set of n basis functions that comprise the columns of an $m \times n$ matrix Φ where $m \gg n$. This structure is captured by the following assumption.

ASSUMPTION 4.

(a) The matrix A_k has the form

$$A_k = \Phi' M_k \Phi,$$

where Φ is an $m \times n$ matrix, and M_k is an $m \times m$ matrix that converges to a matrix M such that $y'My > 0$ for all $y \in \mathbf{R}(\Phi)$ with $y \neq 0$.

(b) The vector b_k has the form

$$b_k = \Phi' d_k, \tag{25}$$

where d_k is a vector in $\Re^m$ that converges to some vector d.

(c) The matrix G_k converges to a matrix G satisfying Assumption 1 together with $A = \Phi' M \Phi$, and satisfies for all k

$$G_k \mathbf{R}(\Phi') \subset \mathbf{R}(\Phi'). \tag{26}$$

We have the following proposition.

PROPOSITION 7. *Let Assumption 4 hold. Then the assumptions of Prop. 6 are satisfied, and there exists $\overline{\gamma} > 0$ such that for all $\gamma \in (0, \overline{\gamma}]$ and all initial iterates x_0, the sequence $\{x_k\}$ generated by iteration (15) converges to a solution of $Ax = b$ with probability 1.*

PROOF. Assumption 4(c) implies Assumption 1, so parts (a) and (b) of Assumption 3 are satisfied. According to the analysis of Prop. (3)(i), Assumption 4(a) implies that

$$\mathbf{N}(\Phi) = \mathbf{N}(A) = \mathbf{N}(A') = \mathbf{N}(A_k) = \mathbf{N}(A'_k) = \mathbf{N}(GA) = \mathbf{N}(G_k A_k). \tag{27}$$

Thus, parts (a) and (c) of Assumption 4 imply Assumption 3, and together with Assumption 4(b), they imply Assumption 2 as well.

From Eq. (27), we have

$$\mathbf{R}(\Phi') = \mathbf{N}(\Phi)^\perp = \mathbf{N}(A)^\perp = \mathbf{N}(A_k)^\perp = \mathbf{N}(A'_k)^\perp = \mathbf{R}(A_k).$$

Hence using the assumption $G_k\mathbf{R}(\Phi') \subset \mathbf{R}(\Phi')$ and the form of b_k given in Eq. (25), we have

$$\mathbf{R}(G_kA_k) = G_k\mathbf{R}(A_k) = G_k\mathbf{R}(\Phi') \subset \mathbf{R}(\Phi') = \mathbf{N}(A)^\perp,$$

and

$$G_kb_k \in G_k\mathbf{R}(\Phi') \subset \mathbf{R}(\Phi') = \mathbf{N}(A)^\perp.$$

Hence the conditions (24) are satisfied. □

We now give a few interesting choices of G_k such that Assumption 4 is satisfied and $\{x_k\}$ converges to a solution of $Ax = b$.

PROPOSITION 8. *Let Assumption 4(a),(b) hold, and let G_k have one of the following forms:*

(i) $G_k = I$.
(ii) $G_k = (\Phi'\Xi_k\Phi + \beta I)^{-1}$, *where Ξ_k converges to a positive definite diagonal matrix and β is any positive scalar.*
(iii) $G_k = (A_k + \beta I)^{-1}$, *where β is any positive scalar.*
(iv) $G_k = (A_k'\Sigma^{-1}A_k + \beta I)^{-1}A_k'\Sigma^{-1}$, *where Σ is any positive definite symmetric matrix.*

Then Assumption 4(c) is satisfied. Moreover, there exists $\overline{\gamma} > 0$ such that for all $\gamma \in (0, \overline{\gamma}]$ and all initial iterates x_0, the sequence $\{x_k\}$ generated by iteration (15) converges to a solution of $Ax = b$ with probability 1.

PROOF. First note that the inverses in parts (ii)-(iv) exist [for case (iii), G_k is invertible for sufficiently large k, since from Assumption 4(a), $x'(A_k + \beta I)x = x'\Phi'M_k\Phi x + \beta\|x\|^2 > 0$ for all $x \neq 0$]. We next verify that G_k converges to a limit G, which together $A = \Phi'M\Phi$ satisfies Assumption 1. For cases (i)-(ii), G_k converges to a symmetric positive definite matrix, so Assumption 1 holds; see Prop. 3(i) of [WB11]. For cases (iii) and (iv), these two iterations are proximal algorithms so Assumption 1 holds; see Prop. 3 (iii)-(iv). We are left to verify the condition $G_k\mathbf{R}(\Phi') \subset \mathbf{R}(\Phi')$ [cf. Eq. (26)].

In cases (i)-(iii), we can show that G_k always takes the form

$$G_k = (\Phi'N_k\Phi + \beta I)^{-1},$$

where N_k is an appropriate matrix and β is a positive scalar [to see this for case (i), we take $N_k = 0$ and $\beta = 1$; for case (ii) we take $N_k = \Xi_k$; for case (iii), recall that $A_k = \Phi'M_k\Phi$ and let $N_k = M_k$]. Let $v \in \mathbf{R}(\Phi')$ and let h be given by $h = G_kv$. Since G_k is invertible, we have

$$v = G_k^{-1}h = (\Phi'N_k\Phi + \beta I)h.$$

Since $v \in \mathbf{R}(\Phi')$, we must have

$$\Phi' N_k \Phi h + \beta h \in \mathbf{R}(\Phi').$$

Note that $\Phi' N_k \Phi h \in \mathbf{R}(\Phi')$, so we must have $\beta h \in \mathbf{R}(\Phi')$. Thus we have shown that $h = G_k v \in \mathbf{R}(\Phi')$ for any $v \in \mathbf{R}(\Phi')$, or equivalently,

$$(\Phi' N_k \Phi + \beta I)^{-1} \mathbf{R}(\Phi') \subset \mathbf{R}(\Phi'). \tag{28}$$

In case (iv), we can write G_k in the form of

$$G_k = (\Phi' N_k \Phi + \beta I)^{-1} A_k' \Sigma^{-1},$$

where $N_k = M_k' \Phi \Sigma^{-1} \Phi' M_k$. For any $v \in \mathbf{R}(\Phi')$, we have

$$h = G_k v = (\Phi' N_k \Phi + \beta I)^{-1} (A_k' \Sigma^{-1} v).$$

Note that $A_k' \Sigma^{-1} v = \Phi' M_k' \Phi \Sigma^{-1} v \in \mathbf{R}(\Phi')$, so by applying Eq. (28) we obtain $h \in \mathbf{R}(\Phi')$. Thus our proof of $G_k \mathbf{R}(\Phi') \subset \mathbf{R}(\Phi')$ is complete for all cases (i)-(iv). □

4. Simulation-based proximal iteration with quadratic regularization. We now consider another special case of the simulation-based iteration

$$x_{k+1} = x_k - \gamma G_k (A_k x_k - b_k), \tag{29}$$

where the residuals converge to 0 naturally. It may be viewed as a proximal iteration, applied to the reformulation of $Ax = b$ as the least squares problem

$$\min_{x \in \Re^n} (Ax - b)' \Sigma^{-1} (Ax - b),$$

or equivalently the optimality condition/linear system

$$A' \Sigma^{-1} A x = A' \Sigma^{-1} b, \tag{30}$$

where Σ is a symmetric positive definite matrix. Generally, proximal iterations are applicable to systems involving a positive semidefinite matrix. By considering instead the system (30), we can bypass this requirement, and apply a proximal algorithm to any linear system $Ax = b$ that has a solution, without A being necessarily positive semidefinite, since the matrix $A' \Sigma^{-1} A$ is positive semidefinite for any A.

Consider the following special choice of the scaling matrix G_k and its limit G:

$$G_k = (A_k'\Sigma^{-1}A_k + \beta I)^{-1}A_k'\Sigma^{-1}, \qquad G = (A'\Sigma^{-1}A + \beta I)^{-1}A'\Sigma^{-1},$$

where β is a positive scalar. We use $\gamma = 1$ and write iteration (29) as

$$x_{k+1} = x_k - (A_k'\Sigma^{-1}A_k + \beta I)^{-1}A_k'\Sigma^{-1}(A_k x_k - b_k), \tag{31}$$

which is equivalent to the sequential minimization

$$x_{k+1} = \operatorname{argmin}_{x\in\Re^n}\left\{\frac{1}{2}(A_k x - b_k)'\Sigma^{-1}(A_k x - b_k) + \frac{\beta}{2}\|x - x_k\|^2\right\}$$

that involves the regularization term $(\beta/2)\|x_k - x\|^2$; see case (iv) of Prop. 3.

It can be shown that this iteration is always nonexpansive, i.e.

$$\|I - G_k A_k\| \le 1, \qquad \forall\ k,$$

due to the use of regularization. However, the convergence analysis is complicated when A is singular, in which case

$$\lim_{k\to\infty} \rho(I - G_k A_k) = \rho(I - GA) = 1.$$

Then the mappings from x_k to x_{k+1} are *not* uniformly contractive, i.e., with a uniform modulus bounded by some $\eta \in (0,1)$ for all sufficiently large k. Thus, to prove convergence of the residuals, we must show that the simulation error accumulates at a relatively slow rate in $\mathbf{N}(A)$. For this reason we need some assumption regarding the rate of convergence of A_k, b_k, and G_k like the following.

ASSUMPTION 5. *The simulation error sequence,*

$$E_k = (A_k - A, b_k - b, G_k - G),$$

viewed as a $(2n^2 + n)$*-vector, satisfies*

$$\limsup_{k\to\infty} k^p\, \mathbf{E}\big\{\|E_k\|^{2p}\big\} < \infty,$$

for some $p > 2$.

Assumption 5 applies to most practical situations that involve Monte Carlo sampling, such as i.i.d. sampling, importance sampling, and Markov chain sampling (see the discussions in Section 2.3). Under natural conditions

(e.g., bounded support, subgaussian tail distribution, etc), these simulation methods satisfy Assumption 5 through forms of the central limit theorem and some concentration inequality arguments.[4] A detailed analysis for various situations where Assumption 5 holds requires dealing with technicalities of the underlying stochastic process, and is beyond our scope. Intuitively, Assumption 5 can be validated for most sampling processes that have good

[4] We will give a brief proof that Assumption 5 holds for the case where $\overline{Z}_k = (A_k, b_k) = \frac{1}{k}\sum_{t=1}^{k}(W_t, v_t)$, and $Z_k = (W_k, v_k)$ are i.i.d. Gaussian random variables with mean $\overline{z} = (A, b)$ and covariance I. By the strong law of large numbers we have $\overline{Z}_k \xrightarrow{a.s.} \overline{z}$. We focus on the error $(G_k - G)$. Define the mapping F as

$$F(\overline{Z}_k) = F\big((A_k, b_k)\big) = G_k = (A_k'\Sigma^{-1}A_k + \beta I)^{-1}A_k'\Sigma^{-1},$$

so that $(G_k - G) = F(\overline{Z}_k) - F(\overline{z})$. By using the differentiability of F (which can be verified by using analysis similar to Konda [Kon02]) and a Taylor series expansion, we have

$$\big\|F(\overline{Z}_k) - F(\overline{z})\big\| \le L\|\overline{Z}_k - \overline{z}\| + L\|\overline{Z}_k - \overline{z}\|^2,$$

for $\overline{Z}_k$ within a sufficiently small neighborhood B of $\overline{z}$, where L is a positive scalar. By using the boundedness of F (which can be verified by showing that the singular values of G_k are bounded), we have for some $M > 0$ and all $\overline{Z}_k$ that

$$\big\|F(\overline{Z}_k) - F(\overline{z})\big\| \le M.$$

Denoting by 1_S the indicator function of an event S, we have for any $p > 2$,

(32)
$$\begin{aligned} k^p\mathbf{E}\left[\|G_k - G\|^{2p}\right] &= k^p\mathbf{E}\left[\|F(\overline{Z}_k) - F(\overline{z})\|^{2p}1_{\{\overline{Z}_k\in B\}}\right] + k^p\mathbf{E}\left[\|F(\overline{Z}_k) - F(\overline{z})\|^{2p}1_{\{\overline{Z}_k\notin B\}}\right] \\ &\le k^p\mathbf{E}\left[\left(L\|\overline{Z}_k - \overline{z}\| + L\|\overline{Z}_k - \overline{z}\|^2\right)^{2p}1_{\{\overline{Z}_k\in B\}}\right] + k^p\mathbf{E}\left[M^{2p}1_{\{\overline{Z}_k\notin B\}}\right] \\ &\le k^pL^{2p}\mathbf{E}\left[\|\overline{Z}_k - \overline{z}\|^{2p}\right] + k^pO\left(\mathbf{E}\left[\|\overline{Z}_k - \overline{z}\|^{2p+1}\right]\right) + k^pM^{2p}\mathbf{P}\left(\overline{Z}_k \notin B\right). \end{aligned}$$

By using the i.i.d. Gaussian assumption regarding Z_k, we obtain that $\overline{Z}_k - \overline{z}$ are zero-mean Gaussians with covariance $\frac{1}{k}I$. From this fact and the properties of Gaussian random variables, we have

$$k^p\mathbf{E}\left[\|\overline{Z}_k - \overline{z}\|^{2p}\right] = const, \qquad \forall\ k \ge 0,$$

and

$$k^pO\left(\mathbf{E}\left[\|\overline{Z}_k - \overline{z}\|^{2p+1}\right]\right) \le O\big(k^p/\sqrt{k^{2p+1}}\big) \to 0, \qquad k^p\mathbf{P}\left(\overline{Z}_k \notin B\right) \le O\left(k^pe^{-k}\right) \to 0,$$

By applying the preceding three relations to Eq. (32), we obtain that $k^p\mathbf{E}\left[\|G_k - G\|^{2p}\right]$ is bounded. Similarly we can prove that $k^p\mathbf{E}\left[\|A_k - A\|^{2p}\right]$ and $k^p\mathbf{E}\left[\|b_k - b\|^{2p}\right]$ are bounded. It follows that for any $p > 0$,

$$\limsup_{k\to\infty} k^p\mathbf{E}\left[\|E_k\|^{2p}\right] \le \limsup_{k\to\infty} k^pO\left(\mathbf{E}\left[\|A_k - A\|^{2p} + \|b_k - b\|^{2p} + \|G_k - G\|^{2p}\right]\right) < \infty.$$

Therefore Assumption 5 is satisfied. This analysis can be easily generalized to sampling processes with subgaussian distributions (e.g., bounded support distributions).

tail properties. In the rare cases where the sampling process may involve heavy tail distributions, it is possible to use increasing numbers of samples between consecutive iterations, to ensure that the estimates converge fast enough and satisfy Assumption 5.

The following proposition is the main result of this section.

PROPOSITION 9 (Convergence of Residual for Proximal Iteration with Quadratic Regularization). *Let Assumptions 2 and 5 hold. Then for all initial iterates x_0, the sequence $\{x_k\}$ generated by iteration (31) satisfies $Ax_k - b \to 0$ and $A_k x_k - b_k \to 0$ with probability 1.*

The proof idea is to first argue that x_k may diverge, but at an expected rate of $O(\log k)$, by virtue of the quadratic regularization. As a result, the accumulated error in $\mathbf{N}(A)$ grows at a rate that is too slow to affect the convergence of the residual. We first show a couple of preliminary lemmas.

LEMMA 1. *For all $x, y \in \Re^n$, $n \times n$ matrix B, and scalar $\beta > 0$, we have*

$$\left\|\beta(B'B + \beta I)^{-1}x + (B'B + \beta I)^{-1}B'y\right\|^2 \le \|x\|^2 + \frac{1}{\beta}\|y\|^2.$$

PROOF. First we consider the simple case when $B = \Lambda$ where Λ is a real diagonal matrix. We define z to be the vector

$$z = \beta(\Lambda^2 + \beta I)^{-1}x + (\Lambda^2 + \beta I)^{-1}\Lambda y.$$

The ith entry of z is $z_i = \frac{\beta x_i + \lambda_i y_i}{\lambda_i^2 + \beta}$, where λ_i is the ith diagonal entry of Λ. We have

$$\begin{aligned} z_i^2 &= \frac{\beta^2 x_i^2 + \lambda_i^2 y_i^2 + 2\beta x_i \lambda_i y_i}{\left(\lambda_i^2 + \beta\right)^2} \\ &\le \frac{(\beta^2 + \beta\lambda_i^2)(x_i^2 + y_i^2/\beta)}{(\lambda_i^2 + \beta)^2} \\ &= \frac{\beta}{\lambda_i^2 + \beta}(x_i^2 + y_i^2/\beta) \\ &\le x_i^2 + \frac{1}{\beta}y_i^2, \end{aligned}$$

where the first inequality uses the fact $2\beta x_i \lambda_i y_i \le \beta\lambda_i^2 x_i^2 + \beta y_i^2$. By summing over i, we obtain

$$\|z\|^2 = \sum_{i=1}^n z_i^2 \le \sum_{i=1}^n x_i^2 + \frac{1}{\beta}\sum_{i=1}^n y_i^2 = \|x\|^2 + \frac{1}{\beta}\|y\|^2, \qquad \forall\, x,\, y \in \Re^n.$$

Thus we have proved that

$$\left\|\beta(\Lambda^2+\beta I)^{-1}x+(\Lambda^2+\beta I)^{-1}\Lambda y\right\|^2 \le \|x\|^2+\frac{1}{\beta}\|y\|^2. \tag{33}$$

Next we consider the general case and the singular value decomposition of B: $B=\bar{U}\Lambda\bar{V}'$, where $\bar{U}$ and $\bar{V}$ are real unitary matrices and Λ is a real diagonal matrix. The vector z is

$$\begin{aligned} z &= \beta(B'B+\beta I)^{-1}x+(B'B+\beta I)^{-1}B'y \\ &= \beta\bar{V}(\Lambda^2+\beta I)^{-1}\bar{V}'x+\bar{V}(\Lambda^2+\beta I)^{-1}\Lambda\bar{U}'y, \end{aligned}$$

or

$$\bar{V}'z=\beta(\Lambda^2+\beta I)^{-1}\bar{V}'x+(\Lambda^2+\beta I)^{-1}\Lambda\bar{U}'y.$$

By applying Eq. (33) to the above relation, we obtain

$$\left\|\bar{V}'z\right\|^2 \le \left\|\bar{V}'x\right\|^2+\frac{1}{\beta}\left\|\bar{U}'y\right\|^2.$$

Since $\bar{U},\bar{V}$ are norm-preserving unitary matrices, we finally obtain $\|z\|^2 \le \|x\|^2+\frac{1}{\beta}\|y\|^2$. □

LEMMA 2. *Under the assumptions of Prop. 9, there exists a positive scalar c such that for all initial iterates x_0, the sequence $\{x_k\}$ generated by iteration (31) satisfies*

$$\mathbf{E}\left\{\|x_k\|^{2p}\right\}^{1/p} \le c\log k,$$

where p is the scalar in Assumption 5.

PROOF. By letting $B=\Sigma^{-1/2}A_k$, we have

$$\begin{aligned} &G_k=\left(B'B+\beta I\right)^{-1}B'\Sigma^{-1/2}, \\ &I-G_kA_k=I-\left(B'B+\beta I\right)^{-1}B'B=\beta\left(B'B+\beta I\right)^{-1}, \end{aligned} \tag{34}$$

where the last equality can be verified by multiplying both sides with $B'B+\beta I$ on the left. Letting x^* be an arbitrary solution of $Ax=b$, we may write iteration (31) as

$$x_{k+1}-x^*=(I-G_kA_k)(x_k-x^*)+G_k(b_k-A_kx^*).$$

or equivalently by using Eq. (34),

$$x_{k+1}-x^*=\beta(B'B+\beta I)^{-1}(x_k-x^*)+(B'B+\beta I)^{-1}B'\overline{e}_k, \tag{35}$$

where we define $\overline{e}_k = \Sigma^{-1/2}(b_k - A_k x^*)$. Applying Lemma 1 to Eq. (35), we have

$$\|x_{k+1} - x^*\|^2 \leq \|x_k - x^*\|^2 + \frac{1}{\beta}\|\overline{e}_k\|^2.$$

We take the pth power of both sides of the above relation and then take expectation, to obtain

$$\begin{aligned}\mathbf{E}\big\{\|x_{k+1} - x^*\|^{2p}\big\}^{1/p} &\leq \mathbf{E}\left\{\left(\|x_k - x^*\|^2 + \frac{1}{\beta}\|\overline{e}_k\|^2\right)^p\right\}^{1/p} \\ &\leq \mathbf{E}\big\{\|x_k - x^*\|^{2p}\big\}^{1/p} + \frac{1}{\beta}\mathbf{E}\big\{\|\overline{e}_k\|^{2p}\big\}^{1/p},\end{aligned}$$

where the last inequality follows from the triangle inequality in the L_p space of random variables with $p > 1$. According to Assumption 5, the sequence $\big\{k\mathbf{E}\big\{\|\overline{e}_k\|^{2p}\big\}^{1/p}\big\}$ is bounded, i.e. $\mathbf{E}\big\{\|\overline{e}_k\|^{2p}\big\}^{1/p} = O(1/k)$. From the preceding inequality, by using induction, we obtain for some positive scalar c

$$\mathbf{E}\big\{\|x_{k+1} - x^*\|^{2p}\big\}^{1/p} \leq \mathbf{E}\big\{\|x_0 - x^*\|^{2p}\big\}^{1/p} + \frac{1}{\beta}\sum_{t=1}^{k}\mathbf{E}\big\{\|\overline{e}_t\|^{2p}\big\}^{1/p} \leq c\log k,$$

where for the last inequality we use the fact $\sum_{t=1}^{k}(1/t) \leq 1 + \log k$. □

Now we are ready to establish the main result on the convergence of the residuals for iteration (31).

Proof of Proposition 9. Let V be an orthonormal basis of $\mathbf{N}(A)^\perp$. We multiply iteration (31) with V' on the left and subtract $V'x^*$ from both sides, yielding

$$\begin{aligned}V'(x_{k+1} - x^*) &= V'x_k - V'x^* - V'G_k(A_k x_k - b_k) \\ &= V'x_k - V'x^* - V'G_k(Ax_k - b) + V'G_k((A - A_k)x_k - b + b_k) \\ &= V'(x_k - x^*) - V'G_k AVV'(x_k - x^*) \\ &\quad + V'G_k((A - A_k)x_k - b + b_k),\end{aligned}$$

where in the last equality we have used $b = Ax^* = AVV'x^*$ and $A = AVV'$ since V is an orthonormal basis of $\mathbf{N}(A)^\perp$. Equivalently by defining $z_k = V'(x_k - x^*)$, we obtain

$$z_{k+1} = V'(I - G_k A)Vz_k + w_k, \tag{36}$$

where w_k is given by

$$w_k = V'G_k\big((A - A_k)x_k - (b - b_k)\big), \tag{37}$$

and $b_k - b \xrightarrow{a.s.} 0$. Note that $\big\|V'(I - G_kA)V\big\| \xrightarrow{a.s.} \big\|V'(I - GA)V\big\|$.

Let us focus on the matrix $V'(I - GA)V$. Using Eq. (34), we have

$$V'(I - GA)V = \beta V'(A'\Sigma^{-1}A + \beta I)^{-1}V = \beta\left(V'A'\Sigma^{-1}AV + \beta I\right)^{-1}.$$

Since V is an orthonormal basis matrix for $\mathbf{N}(A)^{\perp}$, the matrix $V'A'\Sigma^{-1}AV$ is symmetric positive definite. Using the fact $\Sigma^{-1} \succeq \|\Sigma\|^{-1}I$, we have

$$(AV)'\Sigma^{-1}(AV) \succeq (AV)'(\|\Sigma\|^{-1}I)(AV) = \|\Sigma\|^{-1}(AV)'AV \succeq \|\Sigma\|^{-1}\sigma^2(AV)I,$$

where we denote by $\sigma(\cdot)$ the smallest singular value of the given matrix, so we have

$$\sigma\big((AV)'\Sigma^{-1}(AV)\big) \geq \|\Sigma\|^{-1}\sigma^2(AV) > 0.$$

Now by combining the preceding relations, we finally obtain

$$\big\|V'(I - GA)V\big\| \leq \frac{\beta}{\sigma\big((AV)'\Sigma^{-1}(AV)\big) + \beta} \leq \frac{\beta}{\|\Sigma\|^{-1}\sigma^2(AV) + \beta} < 1.$$

In iteration (36) since $V'(I - G_kA)V \xrightarrow{a.s.} V'(I - GA)V$, the matrix $V'(I - G_kA)V$ asymptotically becomes contractive with respect to the Euclidean norm. We are left to show that $(A_k - A)x_k$ in Eq. (37) also converges to 0 with probability 1. By using the Cauchy-Schwartz inequality we obtain

$$\mathbf{E}\big\{\|(A_k - A)x_k\|^p\big\} \leq \sqrt{\mathbf{E}\big\{\|(A_k - A)\|^{2p}\big\}\,\mathbf{E}\big\{\|x_k\|^{2p}\big\}} \leq \bar{c}\frac{(\log k)^{p/2}}{k^{p/2}},$$

where $\bar{c}$ is a positive scalar, and the second inequality uses Lemma 2 and Assumption 5 [i.e. $\mathbf{E}\big\{\|A_k - A\|^{2p}\big\} = O(1/k^p)$].

Using the Markov inequality and fact $p/2 > 1$, we have for any $\epsilon > 0$

$$\sum_{k=1}^{\infty}\mathbf{P}\big(\big\|(A_k - A)x_k\big\| > \epsilon\big) \leq \sum_{k=1}^{\infty}\frac{\mathbf{E}\big\{\|(A_k - A)x_k\|^p\big\}}{\epsilon^p} \leq \frac{\bar{c}}{\epsilon^p}\sum_{k=1}^{\infty}\frac{(\log k)^{p/2}}{k^{p/2}} < \infty,$$

so by applying the Borel-Cantelli lemma, we obtain that $\big\|(A_k - A)x_k\big\| < \epsilon$ for all sufficiently large k with probability 1. Since ϵ can be arbitrarily small, we have $(A_k - A)x_k \xrightarrow{a.s.} 0$. It follows that $w_k \xrightarrow{a.s.} 0$, where w_k is given by Eq. (37). In conclusion, iteration (36) eventually becomes strictly contractive with an additive error $w_k \xrightarrow{a.s.} 0$. It follows that $z_k \xrightarrow{a.s.} 0$ so that $Ax_k - b = AVz_k \xrightarrow{a.s.} 0$. Moreover, we have $A_kx_k - b_k = Ax_k - b + (b - b_k) + (A_k - A)x_k$, so $A_kx_k - b_k \xrightarrow{a.s.} 0$ as well. □

The following example shows that under the assumptions of Prop. 9, the iterate sequence $\{x_k\}$ may diverge with probability 1, even though the residual sequence $\{Ax_k - b\}$ is guaranteed to converge to zero.

EXAMPLE 9 (Divergence of Proximal Iteration with Quadratic Regularization). Let $\beta = 1$, $\Sigma = I$, and

$$A_k = \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & e_{1,k} \end{bmatrix}, \qquad b_k = \begin{bmatrix} 0 \\ e_{2,k} \end{bmatrix}, \qquad x_k = \begin{bmatrix} z_k \\ y_k \end{bmatrix},$$

where $\{e_{1,k}\}$ and $\{e_{2,k}\}$ are approximation errors that converge to 0. Then iteration (31) is equivalent to

$$z_{k+1} = \frac{4}{5} z_k, \qquad y_{k+1} = \frac{1}{1 + e_{1,k}^2} y_k + \frac{e_{1,k} e_{2,k}}{1 + e_{1,k}^2}.$$

For an arbitrary initial iterate $y_0 \in \Re$, we select $\{e_{1,k}\}$ and $\{e_{2,k}\}$ deterministically according to the following rule:

$$e_{1,k} = \frac{1}{\sqrt{k}}, \qquad e_{2,k} = \begin{cases} \frac{3}{\sqrt{k}} & \text{if } y_k < 1 \text{ or } y_{k-1} < y_k \le 2, \\ 0 & \text{if } y_k > 2 \text{ or } y_{k-1} > y_k \ge 1. \end{cases}$$

It can be easily verified that $e_{1,k} \to 0$, $e_{2,k} \to 0$, $e_{1,k} = O(1/\sqrt{k})$, and $e_{2,k} = O(1/\sqrt{k})$, so Assumptions 2 and 5 are satisfied. Clearly $z_k \xrightarrow{a.s.} 0$ so $Ax_k - b \xrightarrow{a.s.} 0$. We will show that the sequence $\{y_k\}$ is divergent.

When $y_k < 1$ or $y_{k-1} < y_k \le 2$, we have

$$y_{k+1} = y_k - \frac{\frac{1}{k}}{1 + \frac{1}{k}} y_k + \frac{\frac{3}{k}}{1 + \frac{1}{k}} \ge y_k + \frac{\frac{1}{k}}{1 + \frac{1}{k}} > y_k,$$

so this iteration will repeat until $y_k > 2$. Moreover, eventually we will have $y_{\bar{k}} > 2$ for some $\bar{k} > k$ since $\sum_{k=0}^{\infty} \frac{\frac{1}{k}}{1+\frac{1}{k}} = \infty$. When $y_k > 2$ or $y_{k-1} > y_k \ge 1$, we have

$$y_{k+1} = y_k - \frac{\frac{1}{k}}{1 + \frac{1}{k}} y_k \le y_k - \frac{\frac{1}{k}}{1 + \frac{1}{k}} < y_k,$$

so this iteration will repeat until $y_k < 1$, and eventually we will have $y_{\bar{k}} < 1$ for some $\bar{k} > k$. Therefore the sequence $\{y_k\}$ crosses the two boundaries of the interval $[1, 2]$ infinitely often, implying that $\{y_k\}$ and $\{x_k\}$ are divergent.

To address the case where the residual sequence $\{Ax_k - b\}$ converges but the iterate sequence $\{x_k\}$ diverges, we may aim to extract out of $\{x_k\}$

the convergent portion, corresponding to $\{V'x_k\}$, which would be simple if $\mathbf{N}(A)$ and $\mathbf{N}(A)^\perp$ are known. This motivates us to estimate the orthogonal projection matrix onto $\mathbf{N}(A)^\perp$ using the sequence $\{A_k\}$. If such an estimate is available, we can extract from $\{x_k\}$ a new sequence of iterates that converges to some solution of $Ax = b$ with probability 1. We will return to this approach and the problem of estimating a projection matrix in Section 6. In what follows, we denote by Π_S the Euclidean projection on a general subspace S.

PROPOSITION 10 (Convergence of Iterates Extracted by Using Π_k). *Let Assumptions 2 and 5 hold, and let $\{\Pi_k\}$ be a sequence of matrices such that*

(38)
$$\Pi_k \xrightarrow{a.s.} \Pi_{\mathbf{N}(A)^\perp} = A'(AA')^\dagger A, \qquad \limsup_{k\to\infty} \mathbf{E}\left\{k^p \left\|\Pi_k - \Pi_{\mathbf{N}(A)^\perp}\right\|^{2p}\right\} < \infty,$$

where $p > 2$ is the scalar in Assumption 5. Let

$$\hat{x}_k = \Pi_k x_k,$$

where x_k is given by iteration (31). Then for all initial iterates x_0, the sequence $\{\hat{x}_k\}$ converges with probability 1 to x^*, the solution of $Ax = b$ that has minimum Euclidean norm.

PROOF. From the proof of Prop. 9, we see that $z_k = V'(x_k - x^*) \xrightarrow{a.s.} 0$ and that $y_k = U'(x_k - x^*)$ may diverge at a rate $O(\log k)$, where y_k and z_k are the components of $x_k - x^*$ in the nullspace decomposition $x_k - x^* = Uy_k + Vz_k$. Since $x^* \in \mathbf{N}(A)^\perp$, we have $\Pi_{\mathbf{N}(A)^\perp} x^* = x^*$. By using this fact we have

$$\begin{aligned}
\hat{x}_k - x^* &= \Pi_k x_k - x^* \\
&= \Pi_{\mathbf{N}(A)^\perp} x_k + (\Pi_k - \Pi_{\mathbf{N}(A)^\perp})x_k - x^* \\
&= \Pi_{\mathbf{N}(A)^\perp}(x_k - x^*) + (\Pi_k - \Pi_{\mathbf{N}(A)^\perp})x_k \\
&= \Pi_{\mathbf{N}(A)^\perp}(Uy_k + Vz_k) + (\Pi_k - \Pi_{\mathbf{N}(A)^\perp})x_k.
\end{aligned}$$

Using the facts that $\Pi_{\mathbf{N}(A)^\perp} U = 0$, $\Pi_{\mathbf{N}(A)^\perp} V = V$, and defining $\tilde{E}_k = \Pi_k - \Pi_{\mathbf{N}(A)^\perp}$, we further obtain

$$\hat{x}_k - x^* = Vz_k + O\big(\|\tilde{E}_k\|\,\|x_k\|\big), \tag{39}$$

By using the Cauchy-Schwartz inequality, together with Lemma 2 and the assumption (38), we have for some $c > 0$ that

$$\mathbf{E}\big\{\|\tilde{E}_k\|^p \|x_k\|^p\big\} \le \sqrt{\mathbf{E}\big\{\|\tilde{E}_k\|^{2p}\big\}\,\mathbf{E}\big\{\|x_k\|^{2p}\big\}} \le c\frac{(\log k)^{p/2}}{k^{p/2}}.$$

Thus for any $\epsilon > 0$, using the Markov inequality and the fact $p/2 > 1$,

$$\sum_{k=1}^{\infty} \mathbf{P}\left(\|\tilde{E}_k\|\|x_k\| > \epsilon\right) \leq \sum_{k=1}^{\infty} \frac{\mathbf{E}\{\|\tilde{E}_k\|^p \|x_k\|^p\}}{\epsilon^p} \leq \frac{c}{\epsilon^p} \sum_{k=1}^{\infty} \frac{(\log k)^{p/2}}{k^{p/2}} < \infty,$$

so by applying the Borel-Cantelli lemma, we obtain $\|\tilde{E}_k\|\|x_k\| \leq \epsilon$ for all sufficiently large k with probability 1. Since $\epsilon > 0$ can be made arbitrarily small, we have $\|\tilde{E}_k\|\|x_k\| \xrightarrow{a.s.} 0$. Finally, we return to Eq. (39) and note that $Vz_k \xrightarrow{a.s.} 0$ (cf. Prop. 9). Thus we have shown that both parts in the right-hand side of Eq. (39) converge to 0. It follows that $\hat{x}_k \xrightarrow{a.s.} x^*$. □

We may also consider a generalization of the proximal iteration (31) that replaces Σ with a sequence of time-varying matrices $\{\Sigma_k\}$, given by

$$x_{k+1} = x_k - \left(A_k' \Sigma_k^{-1} A_k + \beta I\right)^{-1} A_k' \Sigma_k^{-1} (A_k x_k - b_k). \tag{40}$$

We have the following result, which is analogous to the results of Props. 9 and 10.

PROPOSITION 11 (Time-varying Σ_k). *Let Assumptions 2 and 5 hold, and let $\{\Sigma_k\}$ be a sequence of symmetric positive definite matrices satisfying for some $\delta > 0$ that*

$$\Sigma_k \preceq \delta I, \qquad \forall\ k.$$

Then for all initial iterates x_0, the sequence $\{x_k\}$ generated by iteration (40) satisfies $Ax_k - b \to 0$ and $A_k x_k - b_k \to 0$ with probability 1. In addition, if $\{\Pi_k\}$ is a sequence of matrices satisfying the assumptions of Prop. 10, the sequence $\{\hat{x}_k\}$ generated by $\hat{x}_k = \Pi_k x_k$ converges with probability 1 to x^*, the solution of $Ax = b$ that has minimal Euclidean norm.

PROOF. We see that Lemmas 1 and 2 still hold for a sequence of time-varying matrices $\{\Sigma_k\}$. Let $G_k = \left(A_k' \Sigma_k^{-1} A_k + \beta I\right)^{-1} A_k' \Sigma_k^{-1}$. Using an analysis similar to the main proof of Prop. 9, we can show that

$$\left\|\beta V' \left(A' \Sigma_k^{-1} A + \beta I\right)^{-1} V\right\| \leq \frac{\beta}{(1/\delta)\sigma^2(AV) + \beta} < 1, \qquad \forall\ k.$$

It follows that

$$\limsup_{k\to\infty} \left\|V'(I - G_k A)V\right\| < 1.$$

Thus Eq. (36) is still a contraction with additive error decreasing to 0 almost surely. Now we can follow the corresponding steps of Props. 9 and 10, to show that $Ax_k - b \xrightarrow{a.s.} 0$ and $A_k x_k - b_k \xrightarrow{a.s.} 0$, and under the additional assumptions, that $\hat{x}_k = \Pi_k x_k \xrightarrow{a.s.} x^*$. □

5. Examples of almost sure divergence. So far we have focused on proving convergence of the residual and/or the iterate sequences of the stochastic iteration under various assumptions. In this section we will argue reversely. We will present examples of divergence under various conditions and provide the corresponding proofs; see also Example 9 for a divergent case of the quadratic regularization method. These examples justify the need for:

(a) A stabilization scheme in general cases where both the residual and the iterate sequences may diverge.
(b) The use of a projection Π_k to extract a convergent iterate sequence, in special cases when the iterates diverge but the residuals converge to 0.

We will show that the residual sequence $\{Ax_k - b\}$ does not necessarily converge to 0 in general, even if A_k and b_k converge to A and b, respectively. The intuition is that the stochastic iteration is not asymptotically contractive when A is singular, and the simulation errors tend to accumulate in $\mathbf{N}(A)$ at a fast rate, so when transmitted to $\mathbf{N}(A)^\perp$ they cause the residual to diverge with probability 1.

For simplicity, we consider a 2×2 linear system $Ax = b$ where

$$A = \begin{bmatrix} \frac{1}{2} & 0 \\ 0 & 0 \end{bmatrix}, \qquad b = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \qquad \{x \mid Ax = b\} = \mathbf{R} \left\{ \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}.$$

Also in all of the following examples we let $\gamma = 1$. Our analysis remains valid for any value of γ for which the iteration is contractive in $\mathbf{N}(A)^\perp$. By applying the decomposition of Prop. 1 we write $x_k = \begin{bmatrix} z_k \\ y_k \end{bmatrix}$, so the iteration $x_{k+1} = x_k - \gamma G_k(A_k x_k - b_k)$ is equivalent to

$$\begin{bmatrix} z_{k+1} \\ y_{k+1} \end{bmatrix} = (I - G_k A_k) \begin{bmatrix} z_k \\ y_k \end{bmatrix} + G_k b_k. \tag{41}$$

We consider three examples, where the simulation noise enters in the form of four random sequences $\{e_{i,k}\}$, $i = 1, 2, 3, 4$, that have the form

$$e_{i,k} = \frac{1}{k} \sum_{t=1}^{k} w_{i,t} \xrightarrow{a.s.} 0, \qquad i = 1, 2, 3, 4,$$

where $\{w_{i,k}\}$, $i = 1, 2, 3, 4$, are sequences of i.i.d. random variables in $(-1, 1)$ with mean 0 and variance 1.

EXAMPLE 10 (General Projection Iteration). Let $G_k = I$, $b_k = 0$, and

$$A_k = \begin{bmatrix} \frac{1}{2} - e_{1,k} & -e_{2,k} \\ 0 & -e_{3,k} \end{bmatrix}.$$

Then iteration (41) is written as

$$z_{k+1} = \left(\frac{1}{2} + e_{1,k}\right) z_k + e_{2,k} y_k, \qquad y_{k+1} = \left(\prod_{t=1}^{k} (1 + e_{3,t})\right) y_0.$$

We will show that both z_k and y_k diverge with probability 1, so that x_k and even $\hat{x}_k = \Pi_k x_k$ (with $\Pi_k \xrightarrow{a.s.} \Pi_{\mathbf{N}(A)^\perp}$) diverge with probability 1.

EXAMPLE 11 (General Proximal Iteration). Let $G_k = (A_k + I)^{-1}$, $b_k = 0$, and

$$A_k = \begin{bmatrix} \frac{1}{2} + e_{1,k} & e_{2,k} \\ 0 & e_{3,k} \end{bmatrix}.$$

With a straightforward calculation, we can rewrite iteration (41) as

$$z_{k+1} = \frac{1}{3/2 + e_{1,k}} z_k - \frac{e_{2,k}}{(3/2 + e_{1,k})(1 + e_{3,k})} y_k, \qquad y_{k+1} = \left(\prod_{t=1}^{k} \frac{1}{1 + e_{3,t}}\right) y_0.$$

We will show that both z_k and y_k diverge with probability 1, so that x_k and $\hat{x}_k = \Pi_k x_k$ diverge with probability 1.

In Examples 10 and 11 we can show that the z_k-portion of the iteration has the form

$$z_{k+1} = \eta_k z_k + \Theta(e_{2,k}) y_k, \tag{42}$$

where η_k is a random scalar that converges almost surely to some scalar $\eta \in (0, 1)$, and $\Theta(e_{2,k})$ is a Lipschitz function that involves $e_{2,k}$ such that

$$c_1 \|e_{2,k}\| \le \|\Theta(e_{2,k})\| \le c_2 \|e_{2,k}\|, \tag{43}$$

for some $c_1, c_2 > 0$ and all k. According to Eq. (42), the coupling between $\{z_k\}$ and $\{y_k\}$ is induced by $\Theta(e_{2,k}) y_k$, which is equal to 0 in the case with no simulation noise.

EXAMPLE 12 (Nullspace-Consistent Iteration). Let $G_k = I$, $A_k = A$, and $b_k = \begin{bmatrix} 0 \\ e_{4,k} \end{bmatrix}$. Then iteration (41) is equivalent to

$$z_{k+1} = (1/2) z_k, \qquad y_{k+1} = y_0 + \sum_{t=1}^{k} e_{4,t}.$$

This iteration is nullspace-consistent, and $z_k \xrightarrow{a.s.} 0$ so $Ax_k - b \xrightarrow{a.s.} 0$. We will show that y_k diverges with probability 1.

In what follows, we will first prove for all three examples that the sequences $\{y_k\}$ and $\{x_k\}$ diverge with probability 1. We will then focus on Examples 10 and 11, and show that even $\{z_k\}$ and $\{Ax_k - b\}$ diverge with probability 1.

PROPOSITION 12 (Almost Sure Divergence of Iterates). *In Examples 10-12, if* $x_0(2) \neq 0$,

$$\limsup_{k\to\infty} |y_k| = \infty, \qquad \limsup_{k\to\infty} \|x_k\| = \infty, \qquad w.p.1.$$

PROOF. Since $x_0(2) \neq 0$, we have $y_0 \neq 0$. By using the invariance principle for products of partial sums of i.i.d. positive random variables with mean 1 and variance 1 (see Zhang and Huang [ZH07] Theorem 1), we obtain

$$\left(\prod_{t=1}^{k}(1+e_{3,t})\right)^{1/\sqrt{k}} = \left(\frac{1}{k!}\prod_{t=1}^{k}\sum_{i=1}^{t}(1+w_{3,i})\right)^{1/\sqrt{k}} \xrightarrow{i.d.} e^{\sqrt{2}\mathcal{N}(0,1)},$$

and by using the symmetry of the Gaussian distribution, we further obtain

$$\left(\prod_{t=1}^{k}\frac{1}{1+e_{3,t}}\right)^{1/\sqrt{k}} \xrightarrow{i.d.} e^{\sqrt{2}\mathcal{N}(0,1)}.$$

Therefore in both Examples 10 and 11, we obtain

$$y_k^{1/\sqrt{k}} \xrightarrow{i.d.} e^{\sqrt{2}\mathcal{N}(0,1)}, \tag{44}$$

which implies that $\limsup_{k\to\infty} |y_k| = \infty$ with probability 1.

By using the invariance principle for sums of partial sums of i.i.d. random variables ([Kos09] Theorem 2), there exists some $\sigma > 0$ such that

$$\frac{1}{\sqrt{k}}\sum_{t=1}^{k} e_{4,t} = \frac{1}{\sqrt{k}}\sum_{t=1}^{k}\left(\frac{1}{t}\sum_{i=1}^{t} w_{4,i}\right) \xrightarrow{i.d.} \mathcal{N}(0,\sigma^2),$$

which implies that

$$\limsup_{k\to\infty}\left|y_0 + \sum_{t=0}^{k} e_{4,t}\right| = \infty, \qquad w.p.1. \tag{45}$$

Therefore in Example 12 we have $\limsup_{k\to\infty} |y_k| = \infty$ with probability 1. □

The proof of Prop. 12 demonstrates an important cause of divergence of stochastic iterations when the linear system is singular: if a sequence of random variables converges to 1 at a sufficiently slow rate, their infinite product may diverge to infinity.

Now let us focus on Examples 10 and 11, and consider the behavior of $\{z_k\}$ and $\{Ax_k - b\}$. To do so we need to understand the coupling between z_k and y_k, which corresponds to the term $\Theta(e_{2,k})y_k$ in Eq. (42). We address this in the next lemma.

LEMMA 3. *In Examples 10 and 11, if* $x_0(2) \neq 0$, $\limsup_{k\to\infty} |e_{2,k}y_k| = \infty$ *with probability 1.*

PROOF. First we claim that for an arbitrary scalar $m > 0$, the probability $\mathbf{P}\big(|e_{2,k}y_k| \leq m\big)$ decreases to 0 as $k \to \infty$. To prove this, we note that

$$\big\{|e_{2,k}y_k| \leq m\big\} \subset \big\{|e_{2,k}| < 1/k\big\} \cup \big\{|y_k| \leq mk\big\},$$

therefore

$$\mathbf{P}\big(|e_{2,k}y_k| \leq m\big) \leq \mathbf{P}\big(|e_{2,k}| \leq 1/k\big) + \mathbf{P}\big(|y_k| \leq mk\big).$$

As $k \to \infty$, by using the central limit theorem for sample means of i.i.d. random variables, we have $\sqrt{k}e_{2,k} \xrightarrow{i.d.} \mathcal{N}(0,1)$. Using the limit distributions of $\sqrt{k}e_{2,k}$ and $(y_k)^{1/\sqrt{k}}$, we know that

$$\mathbf{P}\big(|e_{2,k}| \leq 1/k\big) = \mathbf{P}\left(\left|\sqrt{k}e_{2,k}\right| \leq 1/\sqrt{k}\right) \downarrow 0, \tag{46}$$

and we can also show that

$$\mathbf{P}\big(|y_k| \leq mk\big) = \mathbf{P}\big(|y_k|^{1/\sqrt{k}} \leq (mk)^{1/\sqrt{k}}\big) \downarrow 0. \tag{47}$$

To see this, we note that $(mk)^{1/\sqrt{k}} \to 0$ and $y_k^{1/\sqrt{k}} \xrightarrow{i.d.} e^{\sqrt{2}\mathcal{N}(0,1)}$ [cf. Eq. (44)], so we have

$$\begin{aligned}
&\mathbf{P}\left(|y_k|^{1/\sqrt{k}} \leq (mk)^{1/\sqrt{k}}\right)\\
&\quad\to \mathbf{P}\left(X \leq 0 \mid X \text{ is distributed according to } e^{\sqrt{2}\mathcal{N}(0,1)}\right)\\
&\quad= 0.
\end{aligned}$$

Equations (46) and (47) prove that $\mathbf{P}\big(|e_{2,k}y_k| \leq m\big) \downarrow 0$.

Now we will show that $\limsup_{k\to\infty} |e_{2,k}y_k| = \infty$ with probability 1. By using the union bound and the continuity of probability measure, we have for any $m > 0$ that

$$\begin{aligned}\mathbf{P}\left(\limsup_{k\to\infty} |e_{2,k}y_k| < m\right) &= \mathbf{P}\left(\bigcup_{n=0}^{\infty}\bigcap_{k\geq n}^{\infty}\{|e_{2,k}y_k| \leq m\}\right)\\ &\leq \sum_{n=0}^{\infty}\lim_{k\to\infty}\mathbf{P}\big(|e_{2,k}y_k| \leq m\big)\\ &= 0,\end{aligned}$$

so that

$$\begin{aligned}\mathbf{P}\left(\limsup_{k\to\infty} |e_{2,k}y_k| < \infty\right) &= \mathbf{P}\left(\bigcup_{m=1}^{\infty}\left\{\sup_{k\geq n}|e_{2,k}y_k| < m\right\}\right)\\ &\leq \sum_{m=1}^{\infty}\mathbf{P}\left(\limsup_{k\to\infty}|e_{2,k}y_k| < m\right)\\ &= 0.\end{aligned}$$

This completes the proof. □

Finally, we are ready to prove the divergence of the sequence $\{z_k\}$ and the residual sequence $\{Ax_k - b\}$ for general projection/proximal iterations based on simulation.

PROPOSITION 13 (Almost Sure Divergence of Residuals). *In Examples 10 and 11, if* $x_0(2) \neq 0$, *then*

$$\limsup_{k\to\infty}\|Ax_k - b\| = \infty, \qquad w.p.1.$$

PROOF. We will focus on the iteration (42) for z_k, i.e.,

$$z_{k+1} = \eta_k z_k + \Theta(e_{2,k})y_k,$$

where $\eta_k \xrightarrow{a.s.} \eta \in (0,1)$ and $\Theta(e_{2,k})$ satisfies Eq. (43). For an arbitrary sample trajectory and sufficiently small $\epsilon > 0$, we have

$$|z_{k+1}| \geq \big|\Theta(e_{2,k})y_k\big| - (\eta + \epsilon)|z_k|,$$

for k sufficiently large. Taking lim sup of both sides of this inequality as $k \to \infty$, and applying Lemma 3, we obtain

$$\limsup_{k\to\infty}|z_{k+1}| \geq \frac{1}{1+\eta+\epsilon}\limsup_{k\to\infty}|\Theta(e_{2,k})y_k| = \infty, \qquad w.p.1.$$

Finally, we have $\limsup_{k\to\infty} \|Ax_k - b\| = \limsup_{k\to\infty} |z_k/2| = \infty$ with probability 1. □

Proposition 13 shows that the residual need not converge to 0 without the nullspace-consistency condition $\mathbf{N}(A) = \mathbf{N}(A_k)$ or some special structure of G_k. The reason is that the simulation error may accumulate in $\mathbf{N}(A)$ through iterative multiplication. This accumulated error in $\mathbf{N}(A)$ corresponds to y_k in our analysis, which diverges at a rate of $e^{\sqrt{k}}$ as proved in Prop. 12. In addition, the simulation error may create a "pathway" from $\mathbf{N}(A)$ to $\mathbf{N}(A)^\perp$ through the mapping from x_k to x_{k+1}, via the term $e_{2,k}$, which decreases to 0 at a rate of $1/\sqrt{k}$. The joint effect is that the accumulated simulation error in $\mathbf{N}(A)$ grows at a rate much faster than the diminishing rate of the "pathway" from $\mathbf{N}(A)$ to $\mathbf{N}(A)^\perp$. As a result, the component z_k corresponding to $\mathbf{N}(A)^\perp$ is "polluted" with simulation error [i.e. $\Theta(e_{2,k})y_k$], which eventually makes the residual diverge with probability 1.

6. Estimating the projection matrix. Finally, we consider the estimation of the matrix of orthogonal projection onto $\mathbf{N}(A)$ or $\mathbf{N}(A)^\perp$, given a sequence of matrices $\{A_k\}$ that converges to A. This is a problem of independent interest, but in our context, the estimated projection matrix can be applied to iterates generated by quadratic regularization [cf. Eq. (31)] and extract a convergent sequence; see Prop. 10 in Section 4. Moreover, this estimate can also be used in the stabilization scheme of selective eigenvalue shifting [cf. Eq. (19)].

Denote the projection matrix to the subspace $\mathbf{N}(A)^\perp$ by

$$\Pi_{\mathbf{N}(A)^\perp} = A'(AA')^\dagger A = VV', \tag{48}$$

where V is an orthonormal basis of $\mathbf{N}(A)^\perp$ and we use $A^\dagger$ to denote the Moore-Penrose pseudoinverse of A. Note that the pseudoinverse is not a continuous operation. The same is true for the decomposition procedure that yields V, since singular vectors are in general unstable if A is perturbed with small error (see Stewart [Ste90] and citations there). The key to constructing a sequence of estimates $\{\Pi_k\}$ based on $\{A_k\}$ such that

$$\Pi_k \xrightarrow{a.s.} \Pi_{\mathbf{N}(A)^\perp},$$

is the convergence of $\mathbf{N}(A_k)$ to $\mathbf{N}(A)$. Although the singular vectors and the inversion of near-zero singular values are sensitive to simulation error, the singular space corresponding to a cluster of singular values is well behaved.

We let the singular value decomposition of A and A_k be

$$A = M\Lambda[U\ V]', \qquad A_k = M_k\Lambda_k[U_k\ V_k]',$$

where Λ and Λ_k are diagonal matrices with diagonal, $\{\lambda_1, \ldots, \lambda_n\}$ and $\{\lambda_{1,k}, \ldots, \lambda_{n,k}\}$, respectively, both ranked in increasing order, and M, M_k, $[U\ V]$, and $[U_k\ V_k]$ are unitary matrices. Assume that the first r singular values of A are equal to 0, and U consists of r basis vectors, so $\mathbf{R}(U) = \mathbf{N}(A)$ and $\mathbf{R}(V) = \mathbf{N}(A)^\perp$. The basis matrices U_k and V_k are chosen so that their column dimensions are equal to those of U and V, respectively.

According to the perturbation theory of singular value decomposition (see the survey [Ste90] and the citations there), the singular values of A_k converge to the singular values of A, and satisfy

$$|\lambda_i - \lambda_{i,k}| = O\big(\|A_k - A\|_F\big), \qquad i = 1, \ldots, n, \qquad w.p.1, \tag{49}$$

where $\|\cdot\|_F$ is the Frobenius matrix norm. Wedin's theorem [Wed12] gives a perturbation bound on a form of angle between the singular subspaces of two matrices assuming some singular value separation conditions. This can also be transformed into a bound on the difference between projection matrices of corresponding singular spaces. A simplified version of Wedin's theorem gives the following bound

$$\|V_k V_k' - VV'\|_F = \|U_k U_k' - UU'\|_F \leq \frac{2\|A - A_k\|_F}{\lambda},$$

where λ is a positive scalar such that

$$\min\{\lambda_{r+1,k}, \ldots, \lambda_{n,k}\} \geq \lambda.$$

We let λ be $\lambda = \lambda_{r+1} - \epsilon$ for some $\epsilon > 0$ sufficiently small. By using the fact $\lambda_{i,k} \xrightarrow{a.s.} \lambda_i$ for each i, it follows that the above condition is satisfied for k sufficiently large, with probability 1. Therefore

$$\|V_k V_k' - VV'\|_F = \|U_k U_k' - UU'\|_F \leq \frac{2\|A - A_k\|_F}{\lambda_{r+1} - \epsilon} = O\big(\|A - A_k\|_F\big), \tag{50}$$

with probability 1; see [SS90]. We will use the Frobenius matrix norm throughout this section, and we note that it is equivalent with the Euclidean matrix norm in the sense that $\|M\| \leq \|M\|_F \leq \sqrt{n}\|M\|$ for any $n \times n$ matrix M.

We now describe an approach for estimating $\Pi_{\mathbf{N}(A)^\perp}$, based on using the singular value decomposition of A_k and applying truncation. In particular, let $\tilde{A}_k$ be obtained by truncating the singular values of A_k that are below a threshold, so near-zero singular values of A_k will be forced to be equal to 0. In order to ensure that only $\{\lambda_{1,k}, \ldots, \lambda_{r,k}\}$ are truncated, we will use a

decreasing sequence of thresholds $\{\delta_k\}$ and assume that it converges to 0 at a rate slower than $\|A_k - A\|_F$. The estimated projection matrix is

$$\Pi_k = \tilde{A}_k' \left(\tilde{A}_k \tilde{A}_k'\right)^\dagger \tilde{A}_k, \tag{51}$$

where $\tilde{A}_k$ is defined by its singular value decomposition

$$\tilde{A}_k = M_k \tilde{\Lambda}_k [U_k \ V_k]', \quad \text{where} \quad \tilde{\Lambda}_k(i,i) = \begin{cases} \lambda_{i,k} & \text{if } \lambda_{i,k} \geq \delta_k, \\ 0 & \text{if } \lambda_{i,k} < \delta_k, \end{cases} \qquad i = 1, \ldots, n.$$

The convergence result of iteration (51) is given in the following proposition.

PROPOSITION 14. *Let $\{\delta_k\}$ be a sequence of positive scalars such that $\delta_k \downarrow 0$ and $(A_k - A)/\delta_k \xrightarrow{a.s.} 0$. Then the sequence $\{\Pi_k\}$ generated by iteration (51) is such that*

$$\Pi_k \xrightarrow{a.s.} \Pi_{\mathbf{N}(A)^\perp},$$

and

$$\left\|\Pi_k - \Pi_{\mathbf{N}(A)^\perp}\right\|_F = O\big(\|A_k - A\|_F\big), \qquad w.p.1.$$

PROOF. We claim that for k sufficiently large, the set of truncated singular values will coincide with the set $\{\lambda_{1,k}, \ldots, \lambda_{r,k}\}$ with probability 1. We first note the almost sure convergence of the singular values of A_k to those of A, i.e.

$$\lambda_{i,k} \xrightarrow{a.s.} 0, \qquad i = 1, \ldots, r, \qquad \lambda_{i,k} \xrightarrow{a.s.} \lambda_i > 0, \qquad i = r+1, \ldots, n.$$

By using Eq. (49) and the assumptions $(A_k - A)/\delta_k \xrightarrow{a.s.} 0$ and $\delta_k \downarrow 0$, we have

$$\begin{aligned} &\lambda_{i,k}/\delta_k = (\lambda_{i,k} - \lambda_i)/\delta_k = O\big(\|A_k - A\|_F\big)/\delta_k \xrightarrow{a.s.} 0, && i = 1, \ldots, r, \\ &\lambda_{i,k} - \delta_k \xrightarrow{a.s.} \lambda_i > 0, && i = r+1, \ldots, n. \end{aligned}$$

This implies that those singular values of A_k whose corresponding singular values of A are equal to 0 will eventually be truncated, and all other singular values will eventually be preserved. Therefore, the truncated diagonal matrix $\tilde{\Lambda}_k$ is equal to

$$\tilde{\Lambda}_k = \operatorname{diag}\{0, \ldots, 0, \lambda_{r+1,k}, \ldots, \lambda_{n,k}\},$$

for k sufficiently large, with probability 1.

Finally, for k sufficiently large, we have

$$\tilde{A}_k = M_k \begin{bmatrix} 0 & 0 \\ 0 & \operatorname{diag}\{\lambda_{r+1,k}, \ldots, \lambda_{n,k}\} \end{bmatrix} \begin{bmatrix} U_k' \\ V_k' \end{bmatrix}, \qquad \mathbf{R}\big(\tilde{A}_k'\big) = \mathbf{R}(V_k), \qquad w.p.1.$$

By the definition of Eq. (51), Π_k is the projection matrix onto $\mathbf{R}\big(\tilde{A}'_k\big)$, so that for all sufficiently large k

$$\Pi_k = V_k V'_k, \qquad w.p.1.$$

This together with Eqs. (48) and (50) yields $\Pi_k \xrightarrow{a.s.} VV' = \Pi_{\mathbf{N}(A)^\perp}$ and $\|\Pi_k - \Pi_{\mathbf{N}(A)^\perp}\|_F = O\big(\|A_k - A\|_F\big)$ with probability 1. □

A second approach for estimating $\Pi_{\mathbf{N}(A)^\perp}$ is to let

$$\Pi_k = A'_k\big(A_k A'_k + \delta_k^2 I\big)^{-1} A_k, \tag{52}$$

where $\{\delta_k\}$ is again a sequence of positive scalars that decreases to 0 at a slow rate.

PROPOSITION 15. *Let $\{\delta_k\}$ satisfy the assumptions of Prop. 14. Then the sequence $\{\Pi_k\}$ generated by iteration (52) is such that*

$$\Pi_k \xrightarrow{a.s.} \Pi_{\mathbf{N}(A)^\perp},$$

and

$$\|\Pi_k - \Pi_{\mathbf{N}(A)^\perp}\|_F = O(\delta_k^2) + O\left(\frac{\|A_k - A\|_F^2}{\delta_k^2}\right), \qquad w.p.1.$$

PROOF. By applying the decomposition $A_k = M_k \Lambda_k [U_k \ V_k]'$ to Eq. (52) we obtain

$$\Pi_k = [U_k \ V_k]\Lambda_k^2(\Lambda_k^2 + \delta_k^2 I)^{-1}[U_k \ V_k]' = [U_k \ V_k]\tilde{I}_k[U_k \ V_k]',$$

where $\tilde{I}_k$ is the diagonal matrix $\tilde{I}_k = \Lambda_k^2\left(\Lambda_k^2 + \delta_k^2 I\right)^{-1}$ that has diagonal elements

$$\tilde{I}_k(i,i) = \frac{\lambda_{i,k}^2}{\lambda_{i,k}^2 + \delta_k^2}, \qquad i = 1, \ldots, n.$$

We also note that

$$\Pi_{\mathbf{N}(A)^\perp} = VV' = [U \ V]\tilde{I}[U \ V]',$$

where $\tilde{I}$ is the diagonal matrix with diagonal elements

$$\tilde{I}(i,i) = \begin{cases} 1 & \text{if } i > r, \\ 0 & \text{if } i \le r, \end{cases} \qquad i = 1, \ldots, n.$$

Let us write

(53)
$$\begin{aligned}\Pi_k - \Pi_{\mathbf{N}(A)^\perp} &= [U_k\ V_k]\tilde{I}_k[U_k\ V_k]' - [U\ V]\tilde{I}[U\ V]' \\ &= [U_k\ V_k]\tilde{I}[U_k\ V_k]' + [U_k\ V_k]\left(\tilde{I}_k - \tilde{I}\right)[U_k\ V_k]' - [U\ V]\tilde{I}[U\ V]' \\ &= V_k V_k' + [U_k\ V_k]\left(\tilde{I}_k - \tilde{I}\right)[U_k\ V_k]' - VV'.\end{aligned}$$

We have

$$\begin{aligned}\left\|\tilde{I} - \tilde{I}_k\right\|_F^2 &= \left\|\tilde{I} - \Lambda_k^2(\Lambda_k^2 + \delta_k^2 I)^{-1}\right\|_F^2 \\ &= \sum_{i=1}^{n}\left|\tilde{I}(i,i) - \frac{\lambda_{i,k}^2}{\lambda_{i,k}^2 + \delta_k^2}\right|^2 \\ &= \sum_{i=1}^{r}\left|\frac{\lambda_{i,k}^2}{\lambda_{i,k}^2 + \delta_k^2}\right|^2 + \sum_{i=r+1}^{n}\left|1 - \frac{\lambda_{i,k}^2}{\lambda_{i,k}^2 + \delta_k^2}\right|^2 \\ &= \sum_{i=1}^{r}\left|\frac{(\lambda_{i,k} - \lambda_i)^2}{(\lambda_{i,k} - \lambda_i)^2 + \delta_k^2}\right|^2 + \sum_{i=r+1}^{n}\left|\frac{\delta_k^2}{\lambda_{i,k}^2 + \delta_k^2}\right|^2 \\ &= O\left(\frac{\|A_k - A\|_F^4}{\delta_k^4}\right) + O\left(\delta_k^4\right),\end{aligned}$$

where the fourth equality uses the fact $\lambda_i = 0$ if $i \le r$, and the last equality uses Eq. (49) to obtain $\lambda_{i,k} - \lambda_i = O\big(\|A_k - A\|_F\big)$ for $i \le r$, and also uses $(\lambda_{i,k} - \lambda_i)/\delta_k \xrightarrow{a.s.} 0$ to obtain $(\lambda_{i,k} - \lambda_i)^2 + \delta_k^2 \approx \delta_k^2$. Therefore

$$\begin{aligned}\left\|\tilde{I} - \tilde{I}_k\right\|_F &\le \sqrt{O\left(\frac{\|A_k - A\|_F^4}{\delta_k^4}\right) + O\left(\delta_k^4\right)} \\ &\le \sqrt{O\left(\frac{\|A_k - A\|_F^4}{\delta_k^4} + \delta_k^4 + 2\|A_k - A\|_F^2\right)} \\ &\le O\left(\frac{\|A_k - A\|_F^2}{\delta_k^2}\right) + O\left(\delta_k^2\right).\end{aligned}$$

By applying the above relation and Eq. (50) to Eq. (53), we obtain

$$\begin{aligned}\|\Pi_k - \Pi_{\mathbf{N}(A)^\perp}\|_F &\le \|V_k V_k' - V'V\|_F + \left\|[U_k\ V_k](\tilde{I}_k - I_k)[U_k\ V_k]'\right\|_F \\ &= \|V_k V_k' - V'V\|_F + \left\|\tilde{I}_k - I_k\right\|_F \\ &= O\big(\|A_k - A\|_F\big) + O(\delta_k^2) + O\left(\frac{\|A_k - A\|_F^2}{\delta_k^2}\right)\end{aligned}$$

$$= O(\delta_k^2) + O\left(\frac{\|A_k - A\|_F^2}{\delta_k^2}\right),$$

where the first equality uses the fact that $\|\cdot\|_F$ is unitarily invariant, and the last equality uses the fact that $2\|A_k - A\|_F \leq \delta_k^2 + \|A_k - A\|_F^2/\delta_k^2$. Finally, since $\delta_k \to 0$ and $\|A_k - A\|_F/\delta_k \to 0$, it follows that $\Pi_k \xrightarrow{a.s.} \Pi_{\mathbf{N}(A)^\perp}$. □

If we let $\delta_k^2 = 1/\sqrt{k}$, which satisfies $(A_k - A)/\delta_k \to 0$ under reasonable assumptions on the convergence rate of $(A_k - A)$, the assumptions of Prop. 15 are satisfied and the bound becomes

$$\|\Pi_k - \Pi_{\mathbf{N}(A)^\perp}\|_F = \frac{1}{\sqrt{k}} O\left(1 + k\|A_k - A\|_F^2\right).$$

Assume that we can verify the assumption (38) of Prop. 10. Then we can apply the estimated projection matrix Π_k to the divergent sequence of iterates $\{x_k\}$ obtained using the quadratic regularization approach. According to Prop. 10, the sequence of iterates $\hat{x}_k = \Pi_k x_k$ converges to the solution of $Ax = b$ with minimal Euclidean norm.

7. Computational illustrations. To illustrate our convergence analysis, we will use the 2×2 problem where

$$A = \begin{bmatrix} 1/2 & 0 \\ 0 & 0 \end{bmatrix}, \qquad b = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \qquad \{x \mid Ax = b\} = \begin{bmatrix} 2 \\ 0 \end{bmatrix} + \mathbf{R}\left\{\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right\},$$

and we will artificially add simulation error to the entries of A and b. In particular, we let A_k and b_k be of the form

$$A_k = A + \frac{1}{k}\sum_{t=1}^{k} W_t, \qquad b_k = b + \frac{1}{k}\sum_{t=1}^{k} w_t, \tag{54}$$

where entries of W_t and w_t are random variables that will be specified later. We will test several algorithms, each with 100 randomly generated iteration trajectories of $\{A_k, b_k\}$, and we will plot, as a function of k, the corresponding "95% confidence interval" of various quantities of interest, which is the range of the 95 values that are closest to the empirical mean value.

EXAMPLE 13 (Estimation of the Stepsize γ). Let A_k and b_k be given by Eq. (54), where all entries of W_k and w_k are i.i.d. Gaussian random variables with mean 0 and variance 1. We generate the sequence $\{\overline{\gamma}_k\}$ using the update rule

$$\overline{\gamma}_{k+1} = \begin{cases} \overline{\gamma}_k & \text{if } \rho(I - \overline{\gamma}_k G_k A_k) \leq 1 + \delta_k, \\ \eta_k \overline{\gamma}_k & \text{if } \rho(I - \overline{\gamma}_k G_k A_k) > 1 + \delta_k, \end{cases}$$

[cf. Section 2.4, Eq. (20)], with

$$G_k = I, \qquad \eta_k = 1 - 1/k, \qquad \delta_k = k^{-1/5}.$$

We start with a randomized $\gamma_0 > 0$ and generate 100 independent trajectories of $\{A_k, b_k, \overline{\gamma}_k\}$.

In Figure 1 we plot the "95% confidence intervals" of the sequences $\{\overline{\gamma}_k\}$ and $\{\rho(I - \overline{\gamma}_k G_k A_k)\}$ respectively. The left side of Figure 1 shows that $\overline{\gamma}_k$ eventually enters the interval $(0, \bar{\gamma})$, where we have $\bar{\gamma} = 4$ for this example. The right side of Figure 1 shows that $\rho(I - \overline{\gamma}_k G_k A_k)$ eventually drops below $1 + \delta_k$, which implies the convergence of $\overline{\gamma}_k$ to an appropriate stepsize.

EXAMPLE 14 (Nullspace-Consistent Iteration for Singular Systems). Let A_k and b_k be given by Eq. (54), where selected entries of W_k and all entries of w_k are i.i.d. Gaussian random variables with mean 0 and variance 0.1. Consider two cases:

(i) *Nullspace-consistent case*: $W_k(1,2) = W_k(2,2) = 0$, while $W_k(1,1)$ and $W_k(2,1)$ are randomly generated as described above.
(ii) *General case*: All entries of W_k are randomly generated.

We generate the sequence $\{x_k\}$ using $x_{k+1} = x_k - \gamma G_k(A_k x_k - b_k)$ with $G_k = I$ and $\gamma = 1$. We start with $x_0 = [10, 10]'$, and we generate 100 independent trajectories of $\{A_k, b_k, x_k\}$.

In Figure 2 we plot the "95% confidence intervals" of the component sequences $\{y_k\}$ and $\{z_k\}$ respectively. The left side of Figure 2 shows that in the nullspace-consistent case (i) the residual sequence converges to 0 and the iterate sequence is unbounded. The right side of Figure 2 shows that in the general case (ii), both the residual and the iterate sequences are unbounded. We have also experimented with the case where γ is estimated using the update rule of Example 13, and have obtained similar results.

EXAMPLE 15 (Proximal Iterations with Quadratic Regularization). Let A_k and b_k be generated as in case (ii) of Example 14, and let $\delta_k = k^{-1/3}$. Consider the following variants of the proximal iteration:

(i) $x_{k+1} = x_k - (A_k' A_k + I)^{-1} A_k'(A_k x_k - b_k)$.
(ii) $x_{k+1} = (1 - \delta_k)x_k - (A_k + I)^{-1}(A_k x_k - b_k)$.
(iii) $x_{k+1} = \big(1 - \delta_k(I - \Pi_k)\big)x_k - (A_k + I)^{-1}(A_k x_k - b_k)$ where Π_k is given by Eq. (52).

Iterations (ii) and (iii) are stabilized [cf. Eqs. (18) and (19)]. We start with $x_0 = [10, 10]'$, and we generate 100 independent trajectories of $\{A_k, b_k, x_k\}$. The "95% confidence intervals" of $\{z_k\}$ are plotted in the left side of Figure 3.

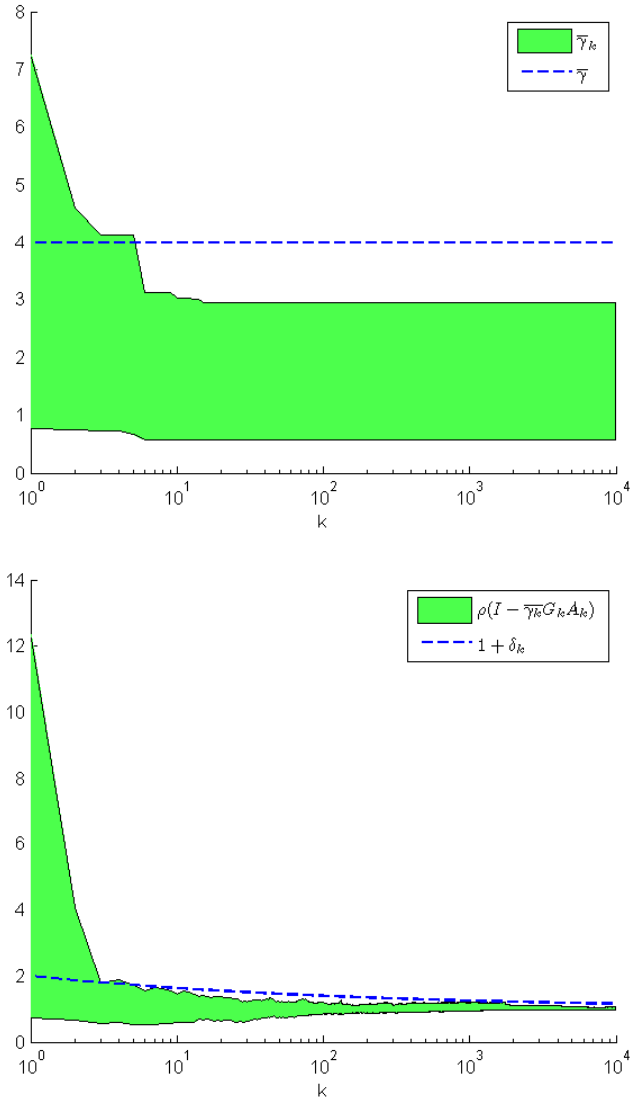

FIG 1. *Estimation of γ using an update rule in Example 13. The top one plots the trajectories of $\overline{\gamma}_k$ against $\bar{\gamma} = 4$, and the bottom one plots the trajectories of $\rho(I - \overline{\gamma}_k G_k A_k)$ against the upperbound $1 + \delta_k$.*

As illustrated by the left side of Figure 3, in iteration (i) where quadratic regularization instead of stabilization is used, the residuals converge to 0 and seem unbiased. By comparison, both stabilized versions of proximal

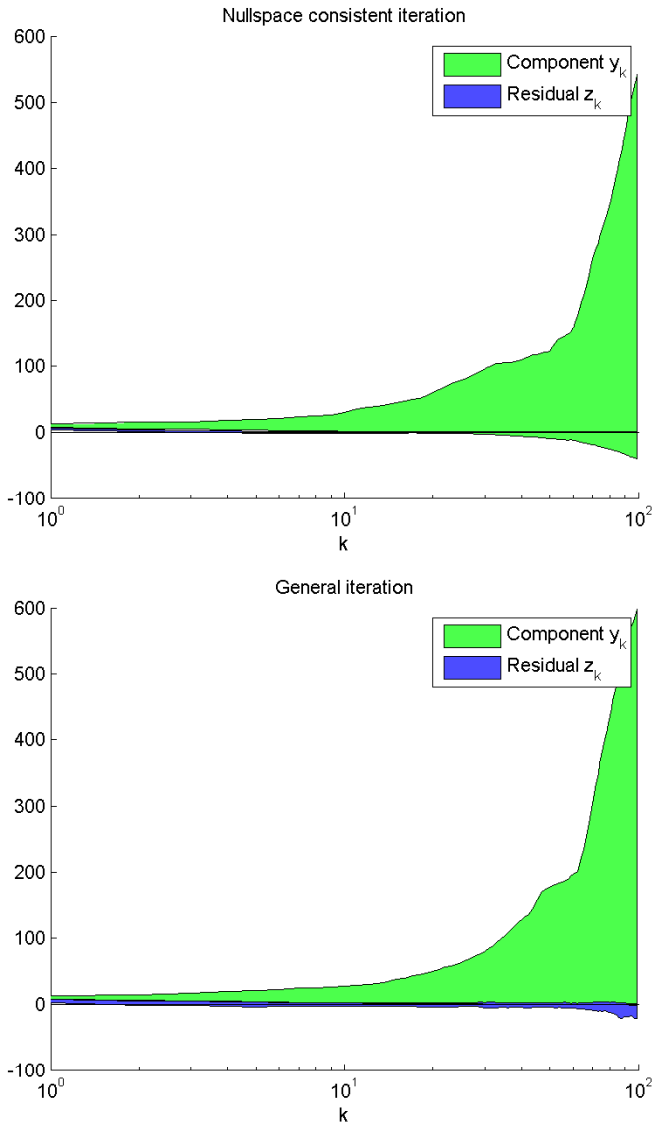

FIG 2. *Convergence of residual of nullspace-consistent iteration in Example 14, compared with the general case. Here the residual coincides with the component z_k. The y_k component of the iterates are divergent in both cases, while the z_k component converges in the nullspace-consistent case and diverges in the general case.*

iteration [(ii) and (iii)] are convergent and biased. Iteration (ii) is subject to decreasing bias in z_k, while iteration (iii) is asymptotically unbiased in z_k since the use of Π_k tends to only perturb y_k.

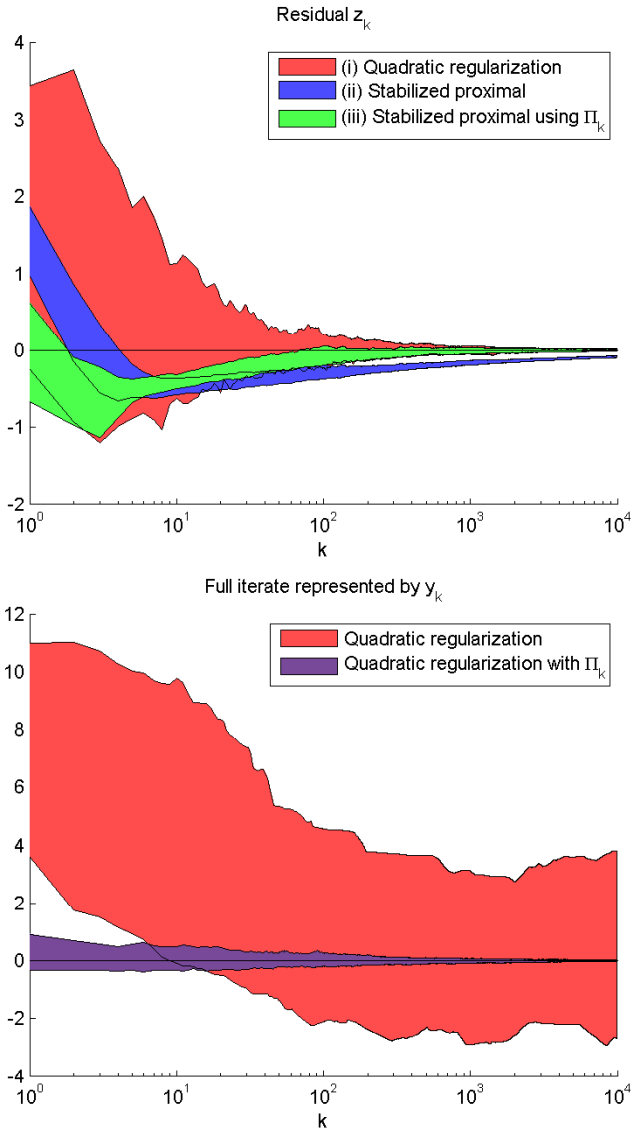

FIG 3. *Convergence of residual for proximal iteration using quadratic regularization in Example 15, compared with stabilized proximal iterations. The top figure shows the residuals z_k of the quadratic regularization algorithm [iteration (i)], compared with those of proximal iterations [(ii) and (iii)] that use stabilization instead. The bottom figure shows the y_k-portion of the iterates x_k generated by iteration (i), and the y_k-portion of the iterates $\hat{x}_k = \Pi_k x_k$. It can be seen that $\{x_k\}$ is divergent, while $\{\hat{x}_k\}$ is convergent.*

TABLE 1
Convergence/Divergence Results of Stochastic Iterative Methods for Singular Systems

Algorithms	**Residuals** $r_k = Ax_k - b$	**Iterates** x_k **and** $\hat{x}_k = \Pi_k x_k$	**References**
General iteration	r_k may diverge.	x_k and $\hat{x}_k$ may diverge.	Examples 10-11
Nullspace-consistent iteration	$r_k \xrightarrow{a.s.} 0$.	x_k may diverge, but $\hat{x}_k$ converges.	Props. 4-7, Example 12
Proximal iteration with quadratic regularization	$r_k \xrightarrow{a.s.} 0$.	x_k may diverge, but $\hat{x}_k$ converges.	Props. 8-10, Example 9

Let $\{x_k\}$ be generated by iteration (i), and let $\{\hat{x}_k\}$ be generated by

$$\hat{x}_k = \Pi_k x_k,$$

where Π_k is given by Eq. (52). The "95% confidence intervals" of the component sequences $\{y_k\}$, are plotted in the right side of Figure 3.

As illustrated by the right side of Figure 3, the sequence generated by quadratic proximal iteration does not converge. This is an example of a stochastic iterative method that generates divergent iterates and convergent residuals. By applying the projection Π_k [estimated using Eq. (52)], we have successfully corrected the divergence of $\{x_k\}$ and extracted a convergent sequence $\{\hat{x}_k\}$.

8. Concluding remarks. We have considered the convergence issues of iterative methods for solving singular linear systems $Ax = b$. Our analysis has focused on simulation-based counterparts of convergent deterministic methods, and has highlighted the complications due to the accumulation of simulation errors along the nullspace of A.

In this paper, we consider two special cases of simulation-based methods: the nullspace-consistent iteration and the propsed proximal algorithm with quadratic regularization. For these cases, we have shown that the residual of the iteration naturally converges to 0, while the iterate may be unbounded. We have categorized through examples and analytical proofs various situations of convergence/divergence of residuals and iterates. To address the issue of divergence, we have proposed a correction method, involving simulation-based estimates of the matrix of projection onto the nullspace of A. A summary of the convergence/divergence results of this paper is given by Table 1.

REFERENCES

[BB96] BRADTKE, S. J. AND BARTO, A. G., Linear least-squares algorithms for temporal difference learning. *Machine Learning*, 22:33–57, 1996.

[Bel70] Bellman, R., *Introduction to Matrix Analysis, 2nd edition*. McGraw-Hill, N. Y, 1970. MR0258847

[Ber10] Bertsekas, D. P., Approximate policy iteration: A survey and some new methods. *Journal of Control Theory and Applications*, 9:310–335, 2010. MR2833999

[Ber11] Bertsekas, D. P., Temporal difference methods for general projected equations. *IEEE Trans. on Automatic Control*, 56:2128–2139, 2011. MR2865769

[Ber12] Bertsekas, D. P., *Dynamic Programming and Optimal Control, Vol. II: Approximate Dynamic Programming*. Athena Scientific, Belmont, M. A, 2012. MR2182753

[BG82] Bertsekas, D. P. and Gafni, E. M., Projection methods for variational inequalities with application to the traffic assignment problem. *Mathematical Programming Study*, 17:139–159, 1982. MR0654697

[BIG74] Ben-Israel, A. and Greville, T. N. E., *Generalized Inverse: Theory and Applications*. Wiley-Interscience, Springer-Verlag, N. Y, 1974. MR0396607

[BMP90] Benveniste, A., Metivier, M., and Priouret, P., *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, NY, 1990. MR1082341

[BN89] Bertsekas, D. P. and Tsitsiklis, J. N. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, Belmont, MA, 1989.

[BN96] Bertsekas, D. P. and Tsitsiklis, J. N. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.

[Bor08] Borkar, V. S., *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, MA, 2008. MR2442439

[Boy02] Boyan, J. A., Technical update: Least-squares temporal difference learning. *Machine Learning*, 49:233–246, 2002.

[BY09] Bertsekas, D. P. and Yu, H., Projected equation methods for approximate solution of large linear systems. *Journal of Computational and Applied Mathematics*, 227:27–50, 2009. MR2512758

[CPS92] Cottle, R. W., Pang, J. S., and Stone, R. E., *The Linear Complementarity Problem*. Academic Press, Boston, MA, 1992. MR1150683

[Dax90] Dax, A., The convergence of linear stationary iterative processes for solving singular unstructured systems of linear equations. *SIAM Review*, 32:611–635, 1990. MR1084572

[DMM06] Drineas, P., Mahoney, M. W., and Muthukrishnan, S., Sampling algorithms for L2 regression and applications. *Proc. 17th Annual SODA*, pages 1127–1136, 2006. MR2373840

[DMMS11] Drineas, P., Mahoney, M. W., Muthukrishnan, S., and Sarlos, T., Faster least squares approximation. *Numerische Mathematik*, 117:219–249, 2011. MR2754850

[FP03] Facchinei, F. and Pang, J. S., *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer, NY, 2003.

[Gol91] Goldberg, J. L., *Matrix Theory with Applications*. McGraw-Hill, N. Y, 1991.

[HY81] Hageman, L. A. and Young, D. M., *Applied Iterative Methods*. Academic Press, NY, 1981. MR0630192

[Kel65] Keller, H. B., On the solution of singular and semidefinite linear systems by iteration. *J. SIAM: Series B, Numerical Analysis*, 2:281–290, 1965. MR0195244

[KNS13] Koshal, J., Nedić, A., and Shanbhag, U. V., Regularized iterative stochastic approximation methods for stochastic variational inequality problems. *IEEE Transactions on Automatic Control*, 58:594–608, 2013.

[Kon02] Konda, V., Actor-critic algorithms. *PhD Thesis, MIT*, 2002.

[Kor76] Korpelevich, G. M., An extragradient method for finding saddle points and for other problems. *Matecon*, 12:747–756, 1976. MR0451121

[Kos09] Kosinski, K. M., On the functional limits for sums of a function of partial sums. *Statistics and Probability Letters*, 79:1552–1527, 2009.

[Kra72] Krasnoselskii, M. A., *Approximate Solution of Operator Equations*. D. Wolters-Noordhoff Pub., Groningen, 1972. MR0385655

[KS13] Kannan, A. and Shanbhag, U. V., Distributed online computation of nash equilibria via iterative regularization techniques. *SIAM Journal of Optimization (to appear)*, 2013.

[KSHdM02] Kleywegt, A. J., Shapiro, A., and Homem-de Mello, T., The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.*, 12:479–502, 2002. MR1885572

[KY03] Kushner, H. J. and Yin, G., *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, NY, 2003. MR1993642

[LP89] Luo, Z. Q. and Tseng, P. On the convergence of a matrix splitting algorithm for the symmetric monotone linear complementarity problem. *SIAM Journal of Control and Optimization*, 29:1037–1060, 1989. MR1110086

[M.70] Sibony, M. Methodes iteratives pour les equations et inequations aux derivees partielles non lineaires de type monotone. *Calcolo*, 7:65–183, 1970. MR0659098

[Mar70] Martinet, B., Regularisation d'inequation variationelles par approximations successives. *Rev. Francaise Inf. Rech. Oper.*, pages 154–159, 1970. MR0298899

[Mey07] Meyn, S. P., *Control Techniques for Complex Systems*. Cambridge University Press, MA, 2007.

[NB03] Nedić, A. and Bertsekas, D. P., Least-squares policy evaluation algorithms with linear function approximation. *Journal of Discrete Event Systems*, 13:79–110, 2003. MR1972051

[NJLS09] Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A., Robust stochastic approximation approach t stochastic programming. *SIAM J. of Optimization*, 19:1574–1609, 2009.

[Pas10] Pasupathy, R., On choosing parameters in retrospective-approximation algorithms for stochastic root finding and simulation optimization. *Operations Research*, 58:889=901, 2010.

[Pow11] Powell, B. W., *Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Ed.* J. Wiley, N. Y, 2011. MR2839330

[Put94] Puterman, M. L., *Markovian Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, New York, NY, 1994. MR1270015

[PWB10] Polydorides, N., Wang, M., and Bertsekas, D. P., A quasi monte carlo method for large-scale linear inverse problems. *Proc. Monte Carlo and Quasi-Monte Carlo, Springer, N. Y*, 2010.

[Roc76] Rockafellar, R. T., Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14:877–898, 1976. MR0410483

[Saa03] Saad, Y., *Iterative Methods for Sparse Linear Systems*. J. SIAM, Philadelphia, P. A, 2003. MR1990645

[SB98] Sutton, R. S. and Barto, A. G., *Reinforcement Learning*. MIT Press, Cambridge, M. A, 1998.

[SDR09] Shapiro, A., Dentcheva, D., and Ruszczynski, A., *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Phila., PA, 2009. MR2562798

[Sha03] Shapiro, A., *Monte Carlo Sampling Methods, in Stochastic Programming*. Handbook in OR & MS, Vol. 10, North-Holland, Amsterdam, 2003. MR2052758

[SS90] Stewart, G. W. and Sun, J. G., *Matrix Perturbation Theory*. Academic Press, Boston, M. A, 1990. MR1061154

[Ste73] Stewart, G. W., *Introduction to Matrix Computations*. Academic Press, NY, 1973. MR0458818

[Ste90] Stewart, G. W., *Perturbation Theory for the Singular Value Decomposition*. CS-TR-2539, College Park, M. D. 1990.

[Tan74] Tanabe, K., Characterization of linear stationary iterative processes for solving a singular system of linear equations. *Numerische Mathematik*, 22:349–359, 1974. MR0351059

[TB97] Trefethen, L. N. and Bau, D., *Numerical Linear Algebra*. J. SIAM, Philadelphia, Philadelphia, P. A, 1997. MR1444820

[WB11] Wang, M. and Bertsekas, D. P., Stabilization of stochastic iterative methods for singular and nearly singular linear systems. *Lab. for Information and Decision Systems Report*, LIDS-P-2878:MIT, Cambridge, MA, 2011.

[Wed12] Wedin, P. A., Perturbation bounds in connection with singular value decomposition. *BIT*, 12:99–111, 2012. MR0309968

[WPB09] Wang, M., Polydorides, N., and Bertsekas, D. P., Approximate simulation-based solution of large-scale least squares problems. *Lab. for Information and Decision Systems Report LIDS-P-2819, MIT, Cambridge, M. A*, 2009.

[YB12] Yu, H. and Bertsekas, D. P., Weighted bellman equations and their applications in dynamic programming. *Lab. for Information and Decision Systems Report LIDS-P-2876, MIT*, 2012.

[You72] Young, D. M., On the consistency of linear stationary iterative methods. *SIAM Journal on Numerical Analysis*, 9:89–96, 1972. MR0312710

[ZH07] Zhang, L. X. and Huang, W., A note on the invariance principle of the product of sums of random variables. *Elect. Comm. in Probability*, 12:51–56, 2007.

Massachusetts Institute of Technology
77 Mass. Ave.
Cambridge, MA 02139
E-mail: mdwang@mit.edu
dimitrib@mit.edu